Coherence in inquisitive first-order logic*

Ivano Ciardelli and Gianluca Grilletti

Preprint — to appear in *Annals of Pure and Applied Logic*

Abstract

Inquisitive first-order logic, InqBQ, is a conservative extension of classical first-order logic with questions. Formulas of InqBQ are interpreted with respect to information states—essentially, sets of relational structures over a common domain. It is unknown whether entailment in InqBQ is compact, and whether validities are recursively enumerable.

In this paper, we study the semantic property of finite coherence: a formula of InqBQ is finitely coherent if in order to determine whether it is satisfied by a state, it suffices to check substates of a fixed finite size.

We show that finite coherence has interesting implications. Most strikingly, entailment towards finitely coherent conclusions is compact.

We identify a broad syntactic fragment of the language, the *rex fragment*, where all formulas are finitely coherent. We give a natural deduction system which is complete for InqBQ entailments with rex conclusions, showing in particular that rex validities are recursively enumerable.

On the way to this result, we study approximations of InqBQ obtained by restricting to information states of a fixed cardinality. We axiomatize the finite approximations and show that, in contrast to the situation in the propositional setting, InqBQ does not coincide with the limit of its finite approximations, settling a question posed by Sano (2011).

Keywords: inquisitive logic, coherence, compactness, axiomatization, team semantics.
Mathematics Subject Classification: 03B45, 03B55, 03B60, 03C07.

*Gefördert durch die Deutsche Forschungsgemeinschaft (DFG), Projektnummer 446711878. Funded by the German Research Foundation (DFG), project number 446711878.

1 Introduction

Inquisitive logic is a research program which aims to extend the scope of logic to questions. In recent years, this extension has been pursued for the language of propositional, first-order, and modal logic, both in the classical and in the non-classical setting.

The standard system of inquisitive first-order logic, InqBQ, can be seen as a conservative extension of classical first-order logic with two new operators that allow us to form questions: inquisitive disjunction, denoted $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$, and the inquisitive existential quantifier, denoted $\exists$ (see Ciardelli, 2009, 2016b; Roelofsen, 2013; Ciardelli *et al.*, 2018; Grilletti, 2020). A defined operator '?' is introduced by letting $?\varphi := \varphi \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} \neg\varphi$. In the language of InqBQ, we have not only standard formulas such as $\forall x Px$, which formalize statements like "every object is P", but also inquisitive formulas such as, for instance, $?\forall x Px$, which formalizes the question "whether every object is P", $\forall x?Px$, which formalizes the question "which objects are P", and $\exists x Px$, which formalizes "what is an instance of P".

Technically, formulas of InqBQ are interpreted not, as usual, in terms of truth relative to a relational structure, but rather in terms of support relative to an *information state*—roughly, a set of relational structures sharing the same domain of quantification.[1] Standard formulas of predicate logic, such as $\forall x Px$, are supported relative to an information state s if they are satisfied point-wise by each structure in s. Inquisitive formulas, by contrast, express *global* properties of information states. Thus, e.g., the formula $?\forall x Px$ expresses the fact that the truth value of $\forall x Px$ is the same in each structure in s; the formula $\forall x?Px$ expresses the fact that the extension of P is the same at each structure in s; and the formula $\exists x Px$ expresses the fact that there is an object d which lies in the extension of P in every structure in s. In InqBQ, it is possible to define a generalized notion of entailment, where formulas standing for questions can occur as premises or conclusions. For example, the valid entailment $\forall x(Px \leftrightarrow \neg Qx), \forall x?Px \models \forall x?Qx$ captures the fact that, given that the extension of P is the complement of the extension of Q, the extension of P determines the extension of Q.

While connectives and quantifiers of InqBQ are well-behaved, satisfying familiar principles, many important problems about the meta-theoretic properties of InqBQ remain open (in spite of much recent work; see Grilletti, 2019, 2020, 2021; Grilletti and Ciardelli, 2021). Most importantly, it is currently unknown if entailment in InqBQ is compact, i.e., if any conclusion that follows from a set of premises also follows from a finite subset of these premises. It is also unknown whether the set of InqBQ-validities is recursively enumerable and, thus, whether

[1] The move taken in inquisitive semantics from single structures to sets of structures is analogous to the move taken in the framework of *team semantics* from single assignments to sets of assignments. Team semantics was first proposed by Hodges (1997a,b), and has been investigated extensively in recent years in connection with several logics of dependence and independence (see, among many others, Väänänen, 2007; Kontinen and Väänänen, 2009; Abramsky and Väänänen, 2009; Galliani, 2012; Grädel and Väänänen, 2013; Yang, 2014, 2019; Yang and Väänänen, 2016). For discussion of the connections between inquisitive and team semantics, see Ciardelli (2016a,b).

a recursive axiomatization exists.

In this paper, we make progress on these key questions by focusing on an important semantic property of formulas: finite coherence. For a cardinal κ, we say that a formula φ is *κ-coherent* if φ is supported by an information state whenever it is supported by all substates of size up to κ. A formula is *finitely coherent* if it is n-coherent for some natural number n. In the team semantics setting, this notion of coherence has been studied by Kontinen (2013), who used it to study the computational complexity of the model checking problem.

Not all formulas of InqBQ are finitely coherent: for instance, $\exists xPx$ is not. On the other hand, many formulas are, and this has a number of important consequences. First, it entails a version of the finite model property: if a finitely coherent formula does not follow from a set of premises, then there is a finite information state that acts as a countermodel (we also show in the paper that this is not the case for arbitrary formulas). Second, finite coherence entails a property known as *normality* in the inquisitive logic literature: if φ is finitely coherent, any information state supporting φ is included in a maximal supporting state. Third, entailment towards finitely coherent conclusions is compact: if a finitely coherent formula follows logically from a set of premises, then it follows from a finite subset of these premises.

We also show that it is possible to identify a syntactic fragment of InqBQ such that all formulas in the fragment are finitely coherent. Since this fragment is obtained by restricting the environments in which the inquisitive existential quantifier $\exists$ is allowed to occur, we will refer to it as the *restricted existential fragment*, or the *rex fragment* for short. The rex fragment is rather broad: it contains all classical formulas, as well as polar questions like $?\forall xPx$ (which ask about the truth value of a statement) and mention-all questions like $\forall x?Px$ (which ask about the extension of a predicate), but not mention-some questions such as $\exists xPx$ (which ask for a witness of a predicate). An interesting question that we will leave open is whether the rex fragment is expressively complete for finitely coherent propositions, i.e., whether any finitely coherent formula in InqBQ is equivalent to one in the fragment.

As we will show, the set of rex validities of InqBQ is recursively enumerable. In fact, we will describe a natural deduction system for InqBQ and show that it is complete with respect to rex conclusions: if a rex formula follows from a set of assumptions, this can be proved in the system. This completeness result is somewhat surprising: all previous known completeness results for InqBQ concern sub-fragments of the *classical antecedent fragment*—the fragment obtained by restricting the antecedents of implications to formulas of standard first-order logic. The most general completeness result known so far, due to Grilletti (2021), covers the entire classical antecedent fragment. The rex fragment is *not* included in the classical antecedent fragment, since it is closed under taking arbitrary implications and thus contains formulas involving inquisitive antecedents, such as $\forall x?Px \to \forall x?Qx$. Thus, our result also shows that allowing questions to freely embed in implications is not by itself an obstacle to completeness, even in the presence of quantification.

Our natural deduction system extends the system for InqBQ described in

Ciardelli (2016b) with a new inference rule, called the *coherence rule*. Essentially, the coherence rule says that in deriving a (verifiably) n-coherent formula, one can freely appeal to the assumption that the relevant state contains at most n possible worlds. This, of course, raises the question of whether the original proof system, without the coherence rule, is itself complete for rex-conclusions. Based on considerations that we will explain in the conclusion, we conjecture that this is not the case, and thus, *a fortiori*, that the system proposed in Ciardelli (2016b) is not complete as a system for InqBQ. However, we leave this as an open question.

In order to reach this result, we make a detour through a question that has independent interest. We consider versions of InqBQ obtained by restricting the class of admissible models by placing a bound on the number of worlds. In this way, we obtain a sequence of logics InqBQ_κ, for κ a cardinal, which we show to converge to our target logic InqBQ. We settle a question posed already by Sano (2011), showing that InqBQ does not coincide with the limit of its finite approximations InqBQ_n for $n \in \mathbb{N}$. We then proceed to give an axiomatization of each finite approximation InqBQ_n and to prove completeness by a canonical model construction involving some novel technical ideas. Completeness for the rex fragment is obtained easily from these results for finite approximations.

The paper is structured as follows. In §2 we cover the relevant preliminaries on InqBQ. In §3 we introduce the notion of coherence that is the focus of the paper and prove some simple results about it. In §4 we define finitary translations from InqBQ to classical two-sorted first-order logic and we use them to show that entailment towards finitely coherent conclusions is compact. In §5 we introduce the rex fragment, show that all formulas in this fragment are finitely coherent, and prove that the set of rex-validities is recursively enumerable. In §6 we show how to write formulas that say that the state of evaluation contains at most n distinct worlds. In §7 we study bounded inquisitive logics InqBQ_κ, obtained by fixing a bound κ to the size of the universe and show that InqBQ is not the limit of the finite-bound logics InqBQ_n for $n \in \mathbb{N}$. In §8 we give complete axiomatizations of the finite-bound inquisitive logic InqBQ_n, and in §9 we use this result to establish a completeness result for the rex fragment of InqBQ. We conclude in §10 by discussing open problems and directions for further work.

2 Preliminaries: inquisitive first-order logic

In this section we introduce the system InqBQ of inquisitive first-order logic and the key facts about it that play a role in the paper. For a more thorough introduction, the reader is referred to Ciardelli (2016b) and Grilletti (2020).

Syntax. As customary, a signature Σ is a set of predicate symbols and function symbols, where each $\sigma \in \Sigma$ is associated with an arity $\mathrm{ar}(\sigma) \geq 0$. We treat identity as a particular binary predicate symbol, denoted $=$, which may or may not belong to Σ. Function symbols of arity 0 are called *constant symbols*. A subset of function symbols are designated as being *rigid*. As a notational convention, we denote rigid function symbols using sans-serif font (e.g., f) and non-rigid function symbols with the default math font (e.g., f). If all function symbols in Σ are rigid, we say that Σ is a function-rigid signature.

Terms of the signature Σ are defined from function symbols and a countably infinite set of variables $\mathsf{Var} = \{x_0, x_1, \dots\}$ in the usual way. A term is said to be *rigid* if it contains only variables and rigid function symbols. We denote rigid terms using sans-serif font (e.g., t) and arbitrary terms with the default math font (e.g., t).

The set $\mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ is given by the following BNF definition, where P is an n-ary predicate symbol from Σ, $t_1, \dots, t_n$ are terms of Σ, and x is a variable.

$$\varphi \ ::= \ P(t_1, \dots, t_n) \mid \bot \mid \varphi \wedge \varphi \mid \varphi \to \varphi \mid \varphi \vee\!\!\!\vee \varphi \mid \forall x \varphi \mid \exists\!\!\!\exists x \varphi$$

Free and bound occurrences of variables are defined in the usual way. A *sentence* is a formula without free variables. As usual, a biconditional operator can be defined by letting $\varphi \leftrightarrow \psi \ := \ (\varphi \to \psi) \wedge (\psi \to \varphi)$.

Formulas that do not contain $\vee\!\!\!\vee$ or $\exists\!\!\!\exists$ are called *classical formulas*, and the set of such formulas is denoted $\mathcal{L}_c(\Sigma)$. If α, β are classical formulas, we define $\neg\alpha := \alpha \to \bot$, $\alpha \vee \beta := \neg(\neg\alpha \wedge \neg\beta)$, and $\exists x \alpha := \neg\forall x \neg\alpha$. In this way, the set of classical formulas can be identified with the standard language of first-order predicate logic in the signature Σ.

The operators $\vee\!\!\!\vee$ and $\exists\!\!\!\exists$ are called *inquisitive disjunction* and *inquisitive existential* respectively and are regarded as question-forming operators. We have an additional defined inquisitive operator ?, obtained by letting $?\varphi := \varphi \vee\!\!\!\vee \neg\varphi$.

We allow ourselves to drop reference to Σ whenever this is unproblematic, writing for instance $\mathcal{L}_{\mathsf{InqBQ}}$ and $\mathcal{L}_c$ for the set of all formulas and classical formulas.

Models. A relational information model is a structure $M = \langle W, D, I \rangle$ where W is a set (the *universe* of M, whose elements are called *possible worlds*), D is a non-empty set (the *domain* of M, whose elements are called *individuals*), and I is a function that associates to each world w a map I_w that assigns to each element of the signature a suitable extension—that is, I_w assigns to each n-ary predicate symbol an n-ary relation over D, and to each n-ary function symbol an n-ary function on D. If $P \in \Sigma$ is a predicate symbol, we write P_w instead of $I_w(P)$ for the extension of P at w, and similarly if f is a function symbol we write f_w for $I_w(f)$. The interpretation map I is subject to two constraints:

1. If f is a function symbol designated as rigid, the interpretation of f is required to be rigid in the sense that for all $w, w' \in W$, $\mathsf{f}_w = \mathsf{f}_{w'}$.
2. If Σ contains the identity predicate $=$, then at each world $I_w(=) \subseteq D \times D$ is required to be a congruence, i.e., an equivalence relation $=_w$ such that, if $\overline{d} = \langle d_1, \dots, d_n \rangle$ and $\overline{d'} = \langle d'_1, \dots, d'_n \rangle$ are any tuples of objects in D such that $d_i =_w d'_i$ for $i \leq n$, we have:
   - for every n-ary predicate $P \in \Sigma$, $\overline{d} \in P_w \iff \overline{d'} \in P_w$;

- for every n-ary function symbol $f \in \Sigma$, $f_w(\overline{d}) =_w f_w(\overline{d'})$.[2]

Intuitively, each world $w \in W$ represents a possible state of affairs. Formally, w is associated with a standard relational structure $\mathcal{M}_w$ for Σ. If Σ does not contain the identity predicate $=$, we can simply define $\mathcal{M}_w$ as the structure $\langle D, I_w \rangle$. If Σ does contain the identity predicate, we define $\mathcal{M}_w$ to be the quotient of $\langle D, I_w \rangle$ modulo $=_w$. Note that since $=_w$ is a congruence, this quotient is well-defined. If two worlds $w, w' \in W$ are associated with the same first-order structure, i.e., $\mathcal{M}_w = \mathcal{M}_{w'}$, then w and w' represent the same state of affairs. In this case, we say that w and w' are *duplicates* and we write $w \approx w'$.[3]

A set $s \subseteq W$ of possible worlds is referred to as an *information state* in M. Intuitively, s models the information that the actual state of affairs corresponds to one of the worlds $w \in s$. A substate $t \subseteq s$ thus represents a body of information that encodes all the information available in s and possibly more. The empty state represents the inconsistent body of information.

Semantics. Let $M = \langle W, D, I \rangle$ be a relational information model. The semantics of InqBQ is given by a relation of *support* between information states $s \subseteq W$ and formulas. As usual, this definition is relativized to an assignment g, which is a map from variables to individuals. Given a term t, the extension of t at world $w \in W$ under g is the individual $[t]^M_{w,g}$ (or simply $[t]^g_w$ if the model is clear from the context) defined inductively in the obvious way. Then we have the following support definition.

Definition 2.1 (Support for InqBQ)**.**
Let $M = \langle W, D, I \rangle$ be a relational information model, $s \subseteq W$, and $g : \mathsf{Var} \to D$.

- $M, s \models_g P(t_1, \dots, t_n) \iff \langle [t_1]^g_w, \dots, [t_n]^g_w \rangle \in P_w$ for all $w \in s$
- $M, s \models_g \bot \iff s = \emptyset$
- $M, s \models_g \varphi \land \psi \iff M, s \models_g \varphi$ and $M, s \models_g \psi$
- $M, s \models_g \varphi \to \psi \iff \forall t \subseteq s : M, t \models_g \varphi$ implies $M, t \models_g \psi$
- $M, s \models_g \varphi \mathbin{\vee\!\!\!\vee} \psi \iff M, s \models_g \varphi$ or $M, s \models_g \psi$
- $M, s \models_g \forall x \varphi \iff M, s \models_{g[x \mapsto d]} \varphi$ for all $d \in D$
- $M, s \models_g \exists x \varphi \iff M, s \models_{g[x \mapsto d]} \varphi$ for some $d \in D$

As usual, the modified assignment $g[x \mapsto d]$ is the assignment that maps x to d and coincides with g on all other variables.

Some more notational conventions: we write $M \models_g \varphi$ instead of $M, W \models_g \varphi$; if $\Phi \subseteq \mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ is a set of formulas, we write $M, s \models_g \Phi$ to mean that $M, s \models_g \varphi$ for all $\varphi \in \Phi$; as usual, for sentences the assignment g is irrelevant and reference to it can be dropped; finally, we allow ourselves to leave the model M implicit when this is harmless, and thus write $s \models_g \varphi$ instead of $M, s \models_g \varphi$.

[2]Note that $I_w(=)$ is not required to be the identity relation. This is because the logic aims to model also situations of uncertainty about whether certain objects are in fact the same. Nevertheless, models where $I_w(w)$ is the identity relation at each world are of course allowed as a special case. For a discussion of the treatment of identity in InqBQ, see Ciardelli (2016b).

[3]Note that in order for w and w' to count as duplicates, the associated structures $\mathcal{M}_w$ and $\mathcal{M}_{w'}$ should be identical, and not merely isomorphic.

The following are three important basic properties of the semantics: the first says that all formulas supported at an information state s remain supported at a stronger information state $t \subseteq s$; the second says that the empty state, which models the state of inconsistent information, supports any formula; the third says that the interpretation of a formula at an information state depends only on the worlds in the state, and not on the rest of the model.

- Persistency: if $M, s \models_g \varphi$ and $t \subseteq s$, then $M, t \models_g \varphi$;
- Empty state property: $M, \emptyset \models_g \varphi$ for all formulas φ;
- Locality: $M, s \models_g \varphi \iff M_{|s}, s \models_g \varphi$
  where $M_{|s} = \langle s, D, I_{|s} \rangle$ is the natural restriction of M to the worlds in s.

Truth and truth-conditionality. A formula of InqBQ is said to be *true* at a possible world w of a model if it is supported by the corresponding singleton state $\{w\}$:

$$M, w \models_g \varphi \overset{\text{def}}{\iff} M, \{w\} \models_g \varphi$$

It is straightforward to prove that for a classical formula α, the truth conditions so defined coincide with the standard ones given by Tarskian semantics. More precisely, if $M = \langle W, D, I \rangle$, we have $M, w \models_g \alpha$ if and only if α is true classically in the relational structure $\mathcal{M}_w$ associated to w.

For some formulas of InqBQ, support at an information state coincides with truth at each world in the state. Such formulas are said to be *truth-conditional*.

Definition 2.2. A formula $\varphi \in \mathcal{L}_{\mathsf{InqBQ}}$ is *truth-conditional* if for every model M, state s, and assignment g we have $M, s \models_g \varphi \iff M, w \models_g \varphi$ for all $w \in s$.

The following proposition says that the truth-conditional formulas of InqBQ are, up to logical equivalence, all and only the classical formulas. Here, two formulas φ, ψ are said to be logically equivalent, notation $\varphi \equiv \psi$, if they have exactly the same semantics, i.e., if for every model M, state s, and assignment g we have $M, s \models_g \varphi \iff M, s \models_g \psi$.

Proposition 2.3. For all formulas $\varphi \in \mathcal{L}_{\mathsf{InqBQ}}$, the following are equivalent:

- φ is truth-conditional;
- $\varphi \equiv \alpha$ for some classical formula α.

In particular, then, all classical formulas are truth-conditional. This guarantees that for such formulas, the above support semantics is essentially equivalent to the standard truth-conditional semantics.

On the other hand, formulas that stand for questions are generally not truth-conditional. As an illustration, here are the support conditions for the three questions used as examples in the introduction.

Example 2.4.

- $s \models \forall x ?Px \iff \forall w, w' \in s : P_w = P_{w'}$

- $s \models ?\forall x Px \iff \forall w, w' \in s : (P_w = D \iff P_{w'} = D)$
- $s \models \exists x Px \iff \exists d \in D\ \forall w \in s : d \in P_w$

In words, $\forall x ? Px$ is supported in a state s if all worlds in s agree on the extension of P; $?\forall x Px$ is supported in s if all worlds in s agree on whether or not the extension of P is the entire domain; and $\exists x Px$ is supported in s if there is an individual d such that all worlds in s agree that d is in the extension of P.

Alternatives and normality. An *alternative* for a formula φ in a model M, relative to an assignment g, is a maximal state supporting φ. The set of such alternatives is denoted $\textsc{Alt}^g_M(\varphi)$:

$$\textsc{Alt}^g_M(\varphi) = \{s \subseteq W \mid s \models_g \varphi \text{ and there is no } t \supset s \text{ such that } t \models_g \varphi\}$$

We say that a formula φ is *normal* if a supporting state for φ can always be extended to an alternative.

Definition 2.5 (Normality)**.** A formula φ is *normal* if for all models M and assignments g, if $M, s \models_g \varphi$ then $s \subseteq a$ for some $a \in \textsc{Alt}^g_M(\varphi)$.

If a formula φ is normal, then its semantics in any model is fully captured by its set of alternatives: indeed, persistence and normalities jointly imply that the supporting states for φ are all and only the states included in some alternative.

In inquisitive propositional logic, all formulas are normal (see Ciardelli and Roelofsen, 2011, Prop. 2.10). As noted in Ciardelli (2009), however, this is not the case for InqBQ. The following proposition gives a counterexample.

Proposition 2.6. The formula $\exists x Px$ is not normal.

Proof. Consider a model M whose domain is the set $\mathbb{N}$ of natural numbers and whose universe is $W = \{w_i \mid i \in \mathbb{N}\}$. Suppose $P_{w_i} = \{n \mid n \geq i\}$. It is easy to check that for any state $s \subseteq W$:

$$s \models \exists x Px \iff s \text{ is finite}$$

Since there is no maximal finite subset of W, there is no alternative for $\exists x Px$ in this model. This is a counterexample to normality, since $\exists x Px$ is supported by some states in M. ☐

Entailment. Entailment in InqBQ is defined as preservation of support:

$$\Phi \models \psi \iff \forall M, s, g : M, s \models_g \Phi \text{ implies } M, s \models_g \psi$$

As usual, if $\Phi = \{\varphi_1, \ldots, \varphi_n\}$ is a finite set we write $\varphi_1, \ldots, \varphi_n \models \psi$ instead of $\{\varphi_1, \ldots, \varphi_n\} \models \psi$; in particular, we write $\models \psi$ instead of $\emptyset \models \psi$ and, if this holds, we say that ψ is *valid*. In restriction to classical formulas, entailment boils down to entailment in classical logic. Thus, classical first-order logic can be seen as a syntactic fragment of InqBQ.

Proposition 2.7 (Conservativity over classical first-order logic)**.**
If $\Gamma \cup \{\alpha\} \subseteq \mathcal{L}_c$ then $\Gamma \models \alpha \iff \Gamma$ entails α in classical first-order logic.

On the other hand, there are many valid entailments which involve questions as premises or conclusions. For instance, we have:

$$\forall x(Px \leftrightarrow \neg Qx), \forall x?Px \models \forall x?Qx$$

The validity of this entailment captures the fact that, under the assumption that the extension of P and the extension of Q are complements of each other, then the extension of P logically determines the extension of Q.

As mentioned in the introduction, many fundamental problems about InqBQ remain open. Most strikingly, it is not known whether InqBQ is entailment-compact, that is, if whenever $\Phi \models \psi$ we also have $\Phi_0 \models \psi$ for some finite subset $\Phi_0 \subseteq \Phi$. Neither is it known whether the set of InqBQ-validities is recursively enumerable and, if so, what a complete axiomatization looks like.

Natural deduction system. A sound, but possibly incomplete, natural deduction system for InqBQ is described in Figure 1. This is an adaptation of the natural deduction system given in Ciardelli (2016b), presented in sequent format for convenience.[4] The introduction and elimination rules for each operator are standard; the only subtlety concerns the fact that a universal quantifier can in general only be eliminated towards a rigid term, and an inquisitive existential can only be introduced from such a term; a universal quantifier can be eliminated towards a non-rigid term only provided the relevant formula is classical. In addition, we have a number of extra principles. First, for classical formulas we have a rule of *reductio ad absurdum*, which reflects the fact that the classical fragment of the logic coincides with classical first-order logic.[5] The two split rules allow us to push a classical antecedent through an inquisitive operator. As discussed in detail in Ciardelli (2016b), these principles capture the fact that non-inquisitive formulas denote specific pieces of information. The constant domains rule encodes the fact that all worlds in a model are assumed to share a common domain of individuals. Finally, the rule KF is related to the fact that negations in InqBQ are always equivalent to some classical formula. For detailed discussion of these inference rules and for an illustration of how they can be combined to give proofs of valid inquisitive entailments, see Ciardelli (2016b).

Note that in restriction to classical formulas, our system includes a complete proof system for classical first-order logic. By the conservativity of InqBQ over classical first-order logic, the system is thus complete with respect to entailments among classical formulas.

[4] Note that what we give is still a natural deduction calculus, and not a sequent calculus, since we give introduction and elimination rules for the logical operators, rather than left and right introduction rules (see §2.1.8 Troelstra and Schwichtenberg, 2000, for discussion).

[5] Equivalently, one could add a rule of double negation elimination for classical formulas, which given $\Theta \vdash \neg\neg\alpha$ allows us to infer $\Theta \vdash \alpha$. This is the choice made in Ciardelli (2016b).

$$\textit{Axiom}\quad \dfrac{}{\varphi \vdash \varphi} \qquad \textit{Weakening}\quad \dfrac{\Theta \vdash \varphi}{\Theta, \Theta' \vdash \varphi} \qquad \textit{Falsum}\quad \dfrac{\Theta \vdash \bot}{\Theta \vdash \varphi}$$

Conjunction

$$\dfrac{\Theta \vdash \varphi \qquad \Theta \vdash \psi}{\Theta \vdash \varphi \wedge \psi} \qquad \dfrac{\Theta \vdash \varphi \wedge \psi}{\Theta \vdash \varphi} \qquad \dfrac{\Theta \vdash \varphi \wedge \psi}{\Theta \vdash \psi}$$

Implication

$$\dfrac{\Theta \vdash \varphi \to \psi}{\Theta, \varphi \vdash \psi} \qquad \dfrac{\Theta \vdash \varphi \qquad \Theta \vdash \varphi \to \psi}{\Theta \vdash \psi}$$

Universal quantifier

$$\dfrac{\Theta \vdash \varphi[y/x]}{\Theta \vdash \forall x \varphi}\, y \notin \mathrm{FV}(\Theta) \qquad \dfrac{\Theta \vdash \forall x \varphi}{\Theta \vdash \varphi[\mathsf{t}/x]} \qquad \dfrac{\Theta \vdash \forall x \alpha}{\Theta \vdash \alpha[t/x]}$$

Inquisitive disjunction

$$\dfrac{\Theta \vdash \varphi}{\Theta \vdash \varphi \vee\!\!\!\vee \psi} \qquad \dfrac{\Theta \vdash \psi}{\Theta \vdash \varphi \vee\!\!\!\vee \psi} \qquad \dfrac{\Theta \vdash \varphi \vee\!\!\!\vee \psi \qquad \Theta, \varphi \vdash \chi \qquad \Theta, \psi \vdash \chi}{\Theta \vdash \chi}$$

Inquisitive existential

$$\dfrac{\Theta \vdash \varphi[\mathsf{t}/x]}{\Theta \vdash \exists\!\!\!\exists x \varphi} \qquad \dfrac{\Theta \vdash \exists\!\!\!\exists x \varphi \qquad \Theta, \varphi[y/x] \vdash \psi}{\Theta \vdash \psi}\, y \notin \mathrm{FV}(\Theta \cup \{\psi\})$$

Identity

$$\dfrac{}{\vdash t = t} \qquad \dfrac{\Theta \vdash \varphi[t/x] \qquad \Theta \vdash t = t'}{\Theta \vdash \varphi[t'/x]}$$

Classical reductio ad absurdum

$$\dfrac{\Theta, \neg\alpha \vdash \bot}{\Theta \vdash \alpha}$$

$$\vee\!\!\!\vee\textit{-split}\quad \dfrac{\Theta \vdash \alpha \to \varphi \vee\!\!\!\vee \psi}{\Theta \vdash (\alpha \to \varphi) \vee\!\!\!\vee (\alpha \to \psi)} \qquad \exists\!\!\!\exists\textit{-split}\quad \dfrac{\Theta \vdash \alpha \to \exists\!\!\!\exists x \varphi}{\Theta \vdash \exists\!\!\!\exists x (\alpha \to \varphi)}\, x \notin \mathrm{FV}(\alpha)$$

$$\textit{Constant domains (CD)}\quad \dfrac{\Theta \vdash \forall x (\varphi \vee\!\!\!\vee \psi)}{\Theta \vdash (\forall x \varphi) \vee\!\!\!\vee \psi}\, x \notin \mathrm{FV}(\psi) \qquad \textit{Classicality of negations (KF)}\quad \dfrac{\Theta \vdash \forall x \neg\neg\varphi}{\Theta \vdash \neg\neg\forall x \varphi}$$

Figure 1: A sound, but possibly incomplete, natural deduction system for InqBQ. In these rules, the variable α ranges over classical formulas, while φ and ψ range over arbitrary formulas and Θ, Θ' over finite sets of arbitrary formulas; t denotes a rigid term, while t, t' denote arbitrary terms, which may but need not be rigid. In all rules, terms substituted for x must be free for x in the relevant formula. We regard the left-hand side of a sequent as a finite set of formulas rather than a sequence; in this way, we do not need explicit structural rules of contraction and exchange.

3 Coherence

In this section we define the coherence properties that are the main focus of this paper and show some consequences of these properties.

Definition 3.1 (κ-coherence)**.** For κ a cardinal, we say that a formula $\varphi \in \mathcal{L}_{\mathsf{InqBQ}}$ is κ-*coherent* if for any model M, state s, and assignment g:

$$s \models_g \varphi \iff \text{for all } t \subseteq s \text{ with } \#t \leq \kappa \text{ we have } t \models_g \varphi$$

where $\#t$ denotes the cardinality of t. We say that φ is *coherent* if it is κ-coherent for some cardinal κ, and *finitely coherent* if it is n-coherent for some natural number n. The set of formulas which are finitely coherent is denoted $\mathcal{L}_{\mathsf{fico}}(\Sigma)$, or simply $\mathcal{L}_{\mathsf{fico}}$ when the signature is clear from the context.

Note that the left-to-right direction in the above definition always holds by persistency, so κ-coherence amounts to the requirement that the converse holds as well, i.e., that support at a state is implied by support at all subsets of size at most κ.

Also, note that truth-conditionality is a special case of n-coherence for $n = 1$. Furthermore, note that if φ is κ-coherent then it is also λ-coherent for all $\lambda \geq \kappa$. This justifies the following definition.

Definition 3.2 (Coherence degree)**.** The coherence degree of a formula φ is the least κ (if it exists) such that φ is κ-coherent. If φ is finitely coherent, the coherence degree is a natural number, denoted d_φ.

Some formulas of InqBQ are not κ-coherent for any cardinal κ (and, as a consequence, have no coherence degree). For these formulas there is no *a priori* bound on the cardinality of the substates we need to consider when checking whether the formula is supported at a state. The simplest example of such a formula is $\exists x Px$.

Proposition 3.3. The formula $\exists x Px$ is not κ-coherent for any κ. That is, for every cardinal κ, there exists a model M and a state s such that $s \not\models \exists x Px$ and for all $t \subseteq s$ with $\#t \leq \kappa$ we have $t \models \exists x Px$.

Proof. Consider an arbitrary cardinal κ, and indicate with κ^+ the cardinal successor of κ. Consider the model $M = \langle W, D, I \rangle$ given by:

- $W = \{w_i \mid i < \kappa^+\}$.
- $D = \{d_j \mid j < \kappa^+\}$.
- $d_j \in P_{w_i} \iff i \neq j$.

We have $M, W \not\models \exists x Px$: indeed, for every element $d_j \in D$ we have have $M, W \not\models P(d_j)$, since $d_j \notin P_{w_j}$. However, given any proper subset $t \subset W$ we have $M, t \models \exists x Px$: to see this, let w_j be a world such that $w_j \notin t$; then for any $w_i \in t$ we have $i \neq j$ and so $d_j \in w_i$, which implies $M, t \models P(d_j)$.

Since the cardinality of W is $\kappa^+ > \kappa$, any subset $t \subseteq W$ with $\#t \leq \kappa$ will be a proper subset of W and thus will support $\exists x Px$. Thus, we have found a state where $\exists x Px$ is not supported, while being supported at all subsets of cardinality up to κ. $\square$

The previous result shows that there are formulas that do not have a coherence degree. On the other hand, as we will see in Section 6, for every finite $n \in \mathbb{N}$ we can produce a formula with coherence degree n. We conjecture that these are the two only possible options for formulas in InqBQ.

Conjecture 3.4 (Dichotomy)**.** The coherence degree of a formula φ is either finite or undefined.

Let us now focus on finitely coherent formulas. An interesting feature of these formulas is that they enjoy a finite model property with respect to the size of the universe W.[6]

Proposition 3.5 (Finite model property for finitely coherent formulas)**.**
Let $n \in \mathbb{N}$. If $\Phi \not\models \psi$ and ψ is n-coherent, the entailment can be falsified in a model M based on a universe W containing at most n worlds.

Proof. Suppose $\Phi \not\models \psi$. Then there are M, s and g such that $M, s \models_g \Phi$ but $M, s \not\models_g \psi$. If ψ is n-coherent, there exists a state $t \subseteq s$ of size at most n such that $M, t \not\models_g \psi$. By persistency, $M, t \models_g \Phi$. Then $M_{|t}$, the restriction of M to t, is a model whose universe is t, and thus contains at most n worlds. By locality we have $M_{|t}, t \models_g \Phi$ but $M_{|t}, t \not\models_g \psi$. $\square$

Moreover, whereas formulas of InqBQ are not in general normal, as we saw, finitely coherent formulas always are.

Proposition 3.6 (Finite coherence implies normality)**.**
If φ is finitely coherent, then it is normal; that is, for every model M, state s and assignment g, if $M, s \models_g \varphi$ then $s \subseteq a$ for some alternative $a \in \textsc{Alt}^g_M(\varphi)$.

Proof. Take an arbitrary model M, information state s, and assignment g such that $M, s \models_g \varphi$. Consider the set S of states containing s and supporting φ:

$$S = \{t \subseteq W \mid s \subseteq t \text{ and } M, t \models_g \varphi\}$$

We want to show that S contains a maximal element.

For this, we first claim that for every non-empty chain $C \subseteq S$ we have $\bigcup C \in S$. Towards a contradiction, suppose this is not the case. Then we have a non-empty chain $C \subseteq S$ such that $\bigcup C \notin S$. Since $\bigcup C$ does include s, we must have $M, \bigcup C \not\models_g \varphi$. Since φ is finitely coherent, there must be a subset $t \subseteq \bigcup C$ of cardinality at most d_φ such that $M, t \not\models_g \varphi$. Since $t \subseteq \bigcup C$, every $w \in t$ is included in some element of the chain, and since t is finite, there must be an element $s' \in C$ of the chain such that $t \subseteq s'$. By persistency, since $M, t \not\models_g \varphi$ we also have $M, s' \not\models_g \varphi$. But this contradicts the hypothesis that $C \subseteq S$.

[6]Obviously, without further assumptions on the signature we cannot hope for a finite model property with respect to the domain D, since the set of 1-coherent formulas already includes all formulas of first-order predicate logic, and we know that some of these formulas can only be falsified over infinite domains D.

We have thus shown that every non-empty chain from S has an upper bound in S. By Zorn's lemma, S contains a maximal element a. This means that a is a maximal extension of s such that $M, a \models_g \varphi$, i.e., $s \subseteq a$ and $a \in \textsc{Alt}^g_M(\varphi)$. $\square$

4 Compactness via finitary translations to FOL

Our next task will be to show that InqBQ entailment towards finitely coherent conclusions is compact. In order to achieve this, we develop a family of maps from the language of InqBQ to the language of classical (two-sorted) first-order logic over a modified signature. These maps allow us to emulate the semantics of InqBQ within standard first-order logic, but only provided a finite upper bound to the size of information states is fixed in advance. This becomes interesting in combination with Proposition 3.5, which guarantees that given an entailment with a finitely coherent conclusion, such a finite bound on the size of the states can indeed be fixed without affecting the validity of the entailment.

Signatures. We associate to a signature Σ a corresponding signature Σ^* over two sorts, $\mathtt{w}$ for worlds and $\mathtt{e}$ for individuals. Σ^* is given as follows:

- For every n-ary predicate symbol $R \in \Sigma$, Σ^* contains a predicate symbol R^* of arity $n+1$ where the first argument is of sort $\mathtt{w}$ and the remaining arguments of sort $\mathtt{e}$.
- For every non-rigid n-ary function symbol $f \in \Sigma$, Σ^* contains a function symbol f^* of arity $n+1$ where the first argument is of sort $\mathtt{w}$ and the remaining arguments as well as the output are of sort $\mathtt{e}$.
- For every rigid n-ary function symbol $\mathsf{f} \in \Sigma$, Σ^* contains a function symbol f^* of arity n where the arguments and the output are of sort $\mathtt{e}$.

We denote by $\mathcal{L}^{\mathsf{FOL}}_{\mathsf{w},\mathsf{e}}(\Sigma^*)$ the language of two-sorted first-order predicate logic over Σ^*. We use w, v for variables of type $\mathtt{w}$ in the latter language, and x, y for variables of type $\mathtt{e}$, which we assume to be the same as the variables of $\mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$.

Models. We associate to a relational information model $M = \langle W, D, I\rangle$ for the signature Σ a two-sorted relational structure $M^* = \langle W, D, I^*\rangle$ for Σ^* where:

- For a predicate symbol R: $I^*(R^*)(w, d_1, \dots, d_n) \iff I_w(R)(d_1, \dots, d_n)$
- For a non-rigid function symbol f: $I^*(f^*)(w, d_1, \dots, d_n) = I_w(f)(d_1, \dots, d_n)$
- For a rigid function symbol f: $I^*(\mathsf{f}^*)(d_1, \dots, d_n) = I_w(\mathsf{f})(d_1, \dots, d_n)$ for an arbitrary $w \in W$

It is easy to check that the map $M \mapsto M^*$ is a bijection between relational information models for Σ and two-sorted relational structures for Σ^*.

Translating terms. Given a term t of $\mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ and a world variable w, we define a corresponding term t_{w} of type e of the language $\mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ inductively as follows:

- if t is a variable x then $t_{\mathsf{w}} = x$
- if $t = f(t^1, \dots, t^n)$ where f is non-rigid then $t_{\mathsf{w}} = f^*(\mathsf{w}, t^1_{\mathsf{w}}, \dots, t^n_{\mathsf{w}})$
- if $t = \mathsf{f}(t^1, \dots, t^n)$ where f is rigid then $t_{\mathsf{w}} = \mathsf{f}^*(t^1_{\mathsf{w}}, \dots, t^n_{\mathsf{w}})$

It is straightforward to check that for any relational information model M, assignment g, and term t of $\mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ we have

$$[t]^M_{w,g} = [t_{\mathsf{w}}]^{M^*}_{g[\mathsf{w}\mapsto w]}$$

where $g[\mathsf{w} \mapsto w]$ is an arbitrary assignment that coincides with g on the variables of type e and maps the variable w to w.

Translating formulas relative to finite states. Let $\mathsf{s} = \{\mathsf{w}_1, \dots, \mathsf{w}_n\}$ be a finite nonempty set of world variables. We define for each $\varphi \in \mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ a formula $\mathrm{tr}_{\mathsf{s}}(\varphi) \in \mathcal{L}^{\mathsf{FOL}}_{\mathsf{w},\mathsf{e}}(\Sigma)$ as follows:

$$\begin{array}{lcl}
\mathrm{tr}_{\mathsf{s}}(R(t^1, \dots, t^k)) & = & R^*(\mathsf{w}_1, t^1_{\mathsf{w}_1}, \dots, t^k_{\mathsf{w}_1}) \wedge \dots \wedge R^*(\mathsf{w}_n, t^1_{\mathsf{w}_n}, \dots, t^k_{\mathsf{w}_n}) \\
\mathrm{tr}_{\mathsf{s}}(\bot) & = & \bot \\
\mathrm{tr}_{\mathsf{s}}(\varphi \wedge \psi) & = & \mathrm{tr}_{\mathsf{s}}(\varphi) \wedge \mathrm{tr}_{\mathsf{s}}(\psi) \\
\mathrm{tr}_{\mathsf{s}}(\varphi \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} \psi) & = & \mathrm{tr}_{\mathsf{s}}(\varphi) \vee \mathrm{tr}_{\mathsf{s}}(\psi) \\
\mathrm{tr}_{\mathsf{s}}(\varphi \to \psi) & = & \bigwedge\{\mathrm{tr}_{\mathsf{s}'}(\varphi) \to \mathrm{tr}_{\mathsf{s}'}(\psi) \mid \emptyset \neq \mathsf{s}' \subseteq \mathsf{s}\} \\
\mathrm{tr}_{\mathsf{s}}(\forall x \varphi) & = & \forall x \, \mathrm{tr}_{\mathsf{s}}(\varphi) \\
\mathrm{tr}_{\mathsf{s}}(\exists x \varphi) & = & \exists x \, \mathrm{tr}_{\mathsf{s}}(\varphi)
\end{array}$$

We spell out one example by way of illustration. We have

$$\begin{array}{ll}
 & \mathrm{tr}_{\mathsf{s}}(\forall x(Px \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} Qx)) \\
= & \forall x(\mathrm{tr}_{\mathsf{s}}(Px) \vee \mathrm{tr}_{\mathsf{s}}(Qx)) \\
= & \forall x((P^*(\mathsf{w}_1, x) \wedge \dots \wedge P^*(\mathsf{w}_n, x)) \vee (Q^*(\mathsf{w}_1, x) \wedge \dots \wedge Q^*(\mathsf{w}_n, x)))
\end{array}$$

The key property of the map tr_{s} is given by the following proposition. The proof is a matter of straightforward case-by-case verification, and is therefore omitted.

Proposition 4.1. Let M be a relational information model, g an assignment, and $s = \{w_1, \dots, w_n\}$ a finite nonempty state. Let $\mathsf{s} = \{\mathsf{w}_1, \dots, \mathsf{w}_n\}$ be a set of n world variables and let $g[\mathsf{s} \mapsto s]$ be any two-sorted assignment that coincides with g on variables of type e and which maps the world variable w_i to w_i for $i = 1, \dots, n$. For any formula $\varphi \in \mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ we have:

$$M, s \models_g \varphi \quad \iff \quad M^* \models_{g[\mathsf{s}\mapsto s]} \mathrm{tr}_{\mathsf{s}}(\varphi)$$

Application to entailments with n-coherent conclusions. The following proposition shows that, although the maps tr_{s} are not in general translations from InqBQ to standard first-order logic, they preserve the validity of entailments whose conclusion is n-coherent for $n \leq \#\mathsf{s}$.

Proposition 4.2. Suppose $\Phi \subseteq \mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ and $\psi \in \mathcal{L}_{\mathsf{fico}}(\Sigma)$. Then we have:

$$\Phi \models_{\mathsf{InqBQ}} \psi \quad \iff \quad \mathrm{tr}_{\mathsf{s}}(\Phi) \models_{\mathsf{FOL}} \mathrm{tr}_{\mathsf{s}}(\psi)$$

where $\models_{\mathsf{FOL}}$ denotes entailment in first-order logic, $\mathsf{s} = \{\mathsf{w}_1, \dots, \mathsf{w}_n\}$ is an arbitrary set of n world variables with $n \geq d_\psi$ and $\mathrm{tr}_{\mathsf{s}}(\Phi) = \{\mathrm{tr}_{\mathsf{s}}(\varphi) \mid \varphi \in \Phi\}$.

Proof. Suppose $\Phi \not\models_{\mathsf{InqBQ}} \psi$ and take $n \geq d_\psi$. By Proposition 3.5, we can find a model M, an assignment g, and a state s of cardinality exactly n (if needed, we may duplicate some worlds) such that $M, s \models_g \Phi$ but $M, s \not\models_g \psi$. By Proposition 4.1 we have $M^* \models_{g[\mathsf{s}\mapsto s]} \mathrm{tr}_{\mathsf{s}}(\Phi)$ but $M^* \not\models_{g[\mathsf{s}\mapsto s]} \mathrm{tr}_{\mathsf{s}}(\psi)$, which shows that $\mathrm{tr}_{\mathsf{s}}(\Phi) \not\models_{\mathsf{FOL}} \mathrm{tr}_{\mathsf{s}}(\psi)$.

For the converse direction, suppose $\mathrm{tr}_{\mathsf{s}}(\Phi) \not\models_{\mathsf{FOL}} \mathrm{tr}_{\mathsf{s}}(\psi)$. This means that there is a two-sorted relational structure M' and an assignment g' such that $M' \models_{g'} \mathrm{tr}_{\mathsf{s}}(\Phi)$ but $M' \not\models_{g'} \mathrm{tr}_{\mathsf{s}}(\psi)$. Now let M be the relational information model such that $M^* = M'$ (which exists since the map $M \mapsto M^*$ is a bijection between relational information models for Σ and two-sorted structures for Σ^*). Let g be the assignment defined by $g(x) = g'(x)$ for every individual variable x, and let $s = \{w_1, \dots, w_n\}$ where $w_i = g'(\mathsf{w}_i)$. By the previous theorem, for any $\chi \in \mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ we have:

$$M, s \models_g \chi \quad \iff \quad M^* \models_{g[\mathsf{s}\mapsto s]} \mathrm{tr}_{\mathsf{s}}(\chi) \quad \iff \quad M' \models_{g'} \mathrm{tr}_{\mathsf{s}}(\chi)$$

where the last biconditional holds because g' and $g[\mathsf{s} \mapsto s]$ coincide on all variables which occur free in $\mathrm{tr}_{\mathsf{s}}(\chi)$. This then implies that $M, s \models_g \Phi$ but $M, s \not\models_g \psi$, which shows that $\Phi \not\models_{\mathsf{InqBQ}} \psi$. $\square$

The existence of this limited translation to standard first-order logic implies that InqBQ-entailment is compact whenever the conclusion if finitely coherent.

Theorem 4.3 (Entailment compactness for finitely coherent conclusions)**.**
If $\Phi \models_{\mathsf{InqBQ}} \psi$ and ψ is finitely coherent, there exists a finite subset $\Phi_0 \subseteq \Phi$ such that $\Phi_0 \models_{\mathsf{InqBQ}} \psi$.

Proof. Suppose $\Phi \models_{\mathsf{InqBQ}} \psi$ and ψ is finitely coherent. By Proposition 4.1, for a suitable choice of the parameter s we have $\mathrm{tr}_{\mathsf{s}}(\Phi) \models_{\mathsf{FOL}} \mathrm{tr}_{\mathsf{s}}(\psi)$. By the compactness of first-order logic, there is a finite subset $\Phi_0 \subseteq \Phi$ such that $\mathrm{tr}_{\mathsf{s}}(\Phi_0) \models_{\mathsf{FOL}} \mathrm{tr}_{\mathsf{s}}(\psi)$. Again by Proposition 4.1, it follows that $\Phi_0 \models_{\mathsf{InqBQ}} \psi$. $\square$

5 The rex fragment

The finitely coherent fragment of InqBQ, $\mathcal{L}_{\mathsf{fico}}$, is defined semantically. In this section, we define a syntactic fragment included in $\mathcal{L}_{\mathsf{fico}}$. This fragment is characterized by a restriction on the occurrences of the inquisitive existential quantifier $\exists$, which is allowed to appear only in the antecedent of an implication.

Definition 5.1. The *restricted existential fragment* of InqBQ, or *rex fragment* for short, is the set $\mathcal{L}_{\mathsf{rex}}(\Sigma)$ (or simply $\mathcal{L}_{\mathsf{rex}}$) of formulas given by the following BNF definition, where p ranges over atomic formulas and φ ranges over the full language $\mathcal{L}_{\mathsf{InqBQ}}$:

$$\chi \ ::= \ p \mid \bot \mid \chi \wedge \chi \mid \varphi \to \chi \mid \chi \vee\!\!\!\vee \chi \mid \forall x \chi$$

We refer to $\chi \in \mathcal{L}_{\mathsf{rex}}$ as a *rex formula.* The crucial property of the fragment is that every rex formula is finitely coherent, and in fact n_χ-coherent for some n_χ which is recursively computable from the syntax of χ. Let us make this precise.

Definition 5.2. To each formula $\chi \in \mathcal{L}_{\mathsf{rex}}$ we associate a natural number n_χ inductively as follows:

- $n_p = 1$ if p is an atomic formula or $\bot$
- $n_{\chi \wedge \xi} = \max(n_\chi, n_\xi)$
- $n_{\varphi \to \chi} = n_\chi$
- $n_{\chi \vee\!\!\!\vee \xi} = n_\chi + n_\xi$
- $n_{\forall x \chi} = n_\chi$

Thus, for instance, the number associated to the mention-all question $\forall x ?Px$ is $n_{\forall x ?Px} = n_{?Px} = n_{Px \vee\!\!\!\vee (Px \to \bot)} = n_{Px} + n_{Px \to \bot} = n_{Px} + n_\bot = 1 + 1 = 2$.

Proposition 5.3. Every rex formula χ is n_χ-coherent.

Proof. Note that a formula χ is n-coherent if and only if it satisfies the following condition for every model M, state s and assignment g:

$$M, s \not\models_g \varphi \quad \Longrightarrow \quad \exists t \subseteq s \, [\, \#t \leq n \text{ and } M, t \not\models_g \varphi \,] \qquad (*)$$

We are going to show that, for every formula $\chi \in \mathcal{L}_{\mathsf{rex}}$, the condition $(*)$ holds for arbitrary M, s and g, and for $n = n_\chi$. The proof consists of an induction on the structure of χ. The base cases and the case for conjunction are trivial and left to the reader. We spell out the remaining three cases below, abbreviating "inductive hypothesis" by IH.

If χ is of the form $\xi \vee\!\!\!\vee \zeta$, for M, s and g as above we have:

$$\begin{aligned} & M, s \not\models_g \xi \vee\!\!\!\vee \zeta \\ \Longrightarrow \quad & M, s \not\models_g \xi \text{ and } M, s \not\models_g \zeta \\ \Longrightarrow \quad & \begin{cases} \exists t_1 \subseteq s \, [\, \#t_1 \leq n_\xi \text{ and } M, t_1 \not\models_g \xi \,] \\ \exists t_2 \subseteq s \, [\, \#t_2 \leq n_\zeta \text{ and } M, t_2 \not\models_g \zeta \,] \end{cases} \quad \text{(by IH)} \end{aligned}$$

If we now let $t = t_1 \cup t_2$, this state witnesses the condition $(*)$ for $n = n_{\xi \vee\!\!\!\vee \zeta}$, since $t \subseteq s$, $\#t \leq \#t_1 + \#t_2 \leq n_\xi + n_\zeta = n_{\xi \vee\!\!\!\vee \zeta}$, and by persistency $M, t \not\models_g \xi \vee\!\!\!\vee \zeta$.

If χ is of the form $\xi \to \zeta$, for M, s and g as above we have:

$$
\begin{array}{lll}
 & M, s \not\models_g \xi \to \zeta & \\
\Longrightarrow & \exists t \subseteq s \left\{ \begin{array}{l} M, t \models_g \xi \\ M, t \not\models_g \zeta \end{array} \right. & \\
\Longrightarrow & \exists t \subseteq s \left\{ \begin{array}{l} M, t \models_g \xi \\ \exists u \subseteq t\, [\, \#u \leq n_\zeta \text{ and } M, u \not\models_g \zeta \,] \end{array} \right. & \text{(by IH)} \\
\Longrightarrow & \exists u \subseteq s \text{ such that } \#u \leq n_\zeta \text{ and } \left\{ \begin{array}{l} M, u \models_g \xi \\ M, u \not\models_g \zeta \end{array} \right. & \text{(by persistency)} \\
\Longrightarrow & \exists u \subseteq s \text{ such that } \#u \leq n_{\xi\to\zeta} \text{ and } M, u \not\models_g \xi \to \zeta & (\text{as } n_{\xi\to\zeta} = n_\zeta)
\end{array}
$$

Thus, the state u witnesses the condition $(*)$ for $n = n_{\xi\to\zeta}$.

If χ is of the form $\forall x\xi$, for M, s and g as above we have:

$$
\begin{array}{lll}
 & M, s \not\models_g \forall x\xi & \\
\Longrightarrow & \exists a \in D \text{ such that } M, s \not\models_{g[x\mapsto a]} \xi & \\
\Longrightarrow & \exists a \in D \text{ such that } \exists t \subseteq s \ \big[\, \#t \leq n_\xi \text{ and } M, t \not\models_{g[x\mapsto a]} \xi \,\big] & \text{(by IH)} \\
\Longrightarrow & \exists t \subseteq s \ \big[\, \#t \leq n_\xi \text{ and } \exists a \in D \text{ such that } M, t \not\models_{g[x\mapsto a]} \xi \,\big] & \\
\Longrightarrow & \exists t \subseteq s \ [\, \#t \leq n_{\forall x\xi} \text{ and } M, t \not\models_g \forall x\xi \,] & (\text{as } n_{\forall x\xi} = n_\xi)
\end{array}
$$

Thus, the state t witnesses condition $(*)$ for $n = n_{\forall x\xi}$. $\square$

Note that the number n_χ is not necessarily equal to the coherence degree d_χ of the formula χ: for instance, we have $n_{Px \vee\!\!\!\vee Px} = n_{Px} + n_{Px} = 2$, but since $Px \vee\!\!\!\vee Px \equiv Px$ we have $d_{Px \vee\!\!\!\vee Px} = d_{Px} = 1$. However, since the coherence degree d_χ is defined as the *least* number n for which χ is n-coherent, the previous proposition guarantees that $n_\chi \geq d_\chi$.

According to this result, the (syntactically defined) rex fragment is included in the (semantically defined) finitely coherent fragment, i.e., we have $\mathcal{L}_{\mathsf{rex}} \subseteq \mathcal{L}_{\mathsf{fico}}$. An interesting question is whether the converse inclusion also holds, modulo logical equivalence. In other words, is any finitely coherent formula equivalent to one in the rex fragment? We leave this question open.

Open question. Is it true that every $\varphi \in \mathcal{L}_{\mathsf{fico}}$ is equivalent to some $\psi \in \mathcal{L}_{\mathsf{rex}}$?

Since rex formulas are finitely coherent, it follows from Theorem 4.3 that entailments with rex conclusions are compact. Moreover, using the results in the previous section we can show that the set of rex validities is recursively enumerable.

Theorem 5.4 (Rex validities are recursively enumerable)**.**
The set $\mathrm{Val}_{\mathsf{rex}} = \{\chi \mid \chi \in \mathcal{L}_{\mathsf{rex}} \text{ and } \models_{\mathsf{InqBQ}} \chi\}$ is recursively enumerable.

Proof. We need to show that there is a method to semi-decide whether a given formula χ belongs to the set $\mathrm{Val}_{\mathsf{rex}}$. This amounts to semi-deciding whether the conjunction

$$\chi \in \mathcal{L}_{\mathsf{rex}} \quad \text{and} \quad \models_{\mathsf{InqBQ}} \chi$$

holds. For this, we proceed as follows. First, we check whether χ is a rex formula. This is a decidable matter: we just need to check if all occurrences of an inquisitive existential quantifier are within the antecedent of a conditional.[7] If χ is not a rex formula, we do not return any output. Otherwise, we need to semi-decide whether χ is valid in InqBQ. For this, we first compute the number n_χ recursively. Then we compute the finitary first-order translation $\mathrm{tr}_{\mathsf{s}}(\chi)$ for s a set of world variables of size n_χ. Since $n_\chi \geq d_\chi$, by Proposition 4.2 we have:

$$\models_{\mathsf{InqBQ}} \chi \iff \models_{\mathsf{FOL}} \mathrm{tr}_{\mathsf{s}}(\chi)$$

Thus, our task reduces to semi-deciding whether $\mathrm{tr}_{\mathsf{s}}(\chi)$ is valid in classical first-order logic. This is possible, since validity in first-order logic is semi-decidable. □

Note that this theorem implies that the set of InqBQ-entailments with a finite number of premises and a rex conclusion is also recursively enumerable. This is because we have:

$$\varphi_1, \ldots, \varphi_n \models_{\mathsf{InqBQ}} \chi \iff \models_{\mathsf{InqBQ}} \varphi_1 \wedge \cdots \wedge \varphi_n \to \chi$$

Thus, semi-deciding whether $\varphi_1, \ldots, \varphi_n \models_{\mathsf{InqBQ}} \chi$ reduces to semi-deciding the validity of the formula $\varphi_1 \wedge \cdots \wedge \varphi_n \to \chi$, which is a rex formula since χ is.

The fact that the set of entailments with rex conclusions is r.e. creates some expectation that it is possible to provide a proof system for InqBQ which is complete with respect to such entailments. As we will see in Section 9, such a proof system can indeed by obtained by supplementing the natural deduction system described in Section 2 with a new inference rule specific to rex conclusions.

6 Cardinality formulas

In this section we show that, for many signatures Σ, we can write for each $n \in \mathbb{N}$ an InqBQ-formula C_n^Σ which says that there are at most n worlds in the state, up to duplicates. The formulas C_n^Σ, which we call *cardinality formulas* for Σ, will play a significant role in the axiomatization results of the next two sections.

Recall that two worlds w, w' in a model M are said to be duplicates ($w \approx w'$) if they are associated with the same relational structure, $\mathcal{M}_w = \mathcal{M}_{w'}$. The *essential cardinality* of a state s, denoted $\#_e s$, is the number of worlds in s, up to equivalence modulo $\approx$:

$$\#_e s = \#(s/\approx)$$

To better convey the idea behind the construction of the cardinality formulas, we first show how to define them in the simple case of the signature $\Sigma = \{P^{(1)}\}$ consisting of just one unary predicate symbol. Once the idea is clear we will show how to extend it to other more general cases.

[7] Note that this step would not go through if we had replaced the condition $\chi \in \mathcal{L}_{\mathsf{rex}}$ with the condition $\chi \in \mathcal{L}_{\mathsf{fico}}$: it is not clear whether there is a method to decide, or even semi-decide, whether a given formula χ is finitely coherent.

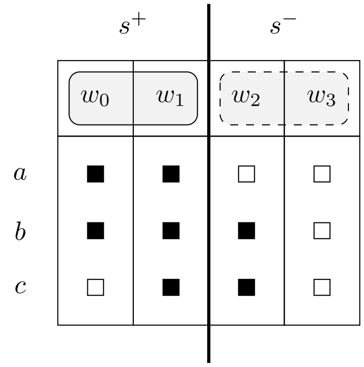

Figure 2: A model M in the signature $\Sigma = \{P^{(1)}\}$ with four worlds and without duplicates. Each row corresponds to an individual and each column to a world. The square in the cell corresponding to individual d and world w is black if $d \in P_w$, and white otherwise. To show that the formula C_4 holds at the model, we can follow the steps in the proof of Proposition 6.2: We can use the element a to partition the worlds of the model, since in the worlds of s^+ it satisfies predicate P and in the worlds of s^- it does not satisfy P. Both s^+ and s^- contain two worlds, thus we have $s^+ \models C_2$ and $s^- \models C_2$. From this it follows that $M \models_{[x \mapsto a]} (P(x) \to C_2) \wedge (\neg P(x) \to C_2)$, which immediately implies $M \models C_4$.

Definition 6.1 (Cardinality formulas, case $\Sigma = \{P^{(1)}\}$)**.**
We define the cardinality formulas $C_n^{\{P\}}$ inductively as follows:

$$C_0^{\{P\}} := \bot$$
$$C_1^{\{P\}} := \forall x ? Px$$
$$C_{n+1}^{\{P\}} := \exists x \bigvee\nolimits_{i=1}^{n} \Big[\, (Px \to C_i^{\{P\}}) \wedge (\neg Px \to C_{n+1-i}^{\{P\}}) \, \Big]$$

Proposition 6.2. Let M be a model for the signature $\Sigma = \{P^{(1)}\}$, s an information state in M and $n \in \mathbb{N}$. We have:

$$M, s \models C_n^{\{P\}} \quad \iff \quad \#_e s \leq n$$

Proof. We may assume without loss of generality that s does not contain duplicate worlds, so that the essential cardinality of s is just its cardinality, $\#_e s = \#s$. If this is not the case, we can instead work with a substate $s' \subseteq s$ obtained by choosing a single representative for each equivalence class modulo $\approx$ which is represented in s. Then s' does not contain duplicate worlds, has the same essential cardinality as s, and it is easy to see that s and s' support the same formulas. Thus, if the claim can be shown for s', it applies to s as well.

We prove the statement by induction on the number n. To lighten the notation, throughout the proof we omit superscripts and just write C_n for $C_n^{\{P\}}$.

The case for $n = 0$ is trivial. For the case $n = 1$, consider the formula C_1. As we have seen in Example 2.4, we have

$$s \models C_1 \iff s \models \forall x?Px \iff \forall w, w' \in s : P_w = P_{w'}$$

Since we are assuming that Σ contains only the predicate P, the above condition says that all worlds $w, w' \in s$, w and w' are duplicates of each other, i.e., that the essential cardinality of s is at most 1.

For the inductive step, assume the property holds for the formulas C_k with $k < n$ and consider the formula C_n for $n \geq 2$. In case $s = \emptyset$, it trivially supports the formula C_n by the empty state property. And in case s contains exactly one world, by inductive hypothesis s supports all formulas of the form C_k for $1 \leq k < n$, and it is easy to see that this implies that s supports C_n too.

Now suppose s contains at least two worlds and at most n. Let w_0 and w_1 be two distinct worlds in s. Since these worlds are not duplicates and P is the only symbol in the language, $P_{w_0} \neq P_{w_1}$, and so there is an element a such that either $a \in (P_{w_0} - P_{w_1})$, or $a \in (P_{w_1} - P_{w_0})$. Without loss of generality, suppose the former. Now define $s^+ = \{w \in s \mid a \in P_w\}$ and $s^- = \{w \in s \mid a \notin P_w\}$. Note that s^+ and s^- are both non-empty (since they contain respectively the worlds w_0 and w_1) and they form a partition of s. The idea is illustrated in Figure 2. For $k := \#s^+$ we have $1 \leq k \leq n-1$, and $s^+ \models C_k$ by inductive hypothesis. Now take any $t \subseteq s$ such that $t \models_{[x \mapsto a]} P(x)$. By the semantics of atoms, this implies $t \subseteq s^+$ and so by persistency also $t \models C_k$. This shows that $s \models_{[x \mapsto a]} P(x) \to C_k$. Moreover $\#s^- = \#s - \#s^+ \leq n-k$, and so by inductive hypothesis $s^- \models C_{n-k}$. Now take any $t \subseteq s$ such that $t \models_{[x \mapsto a]} \neg P(x)$. By the semantics, $t \subseteq s^-$, and so by persistency $t \models C_{n-k}$. Thus, $s \models_{[x \mapsto a]} \neg P(x) \to C_{n-k}$. These two conditions, together with the fact that $1 \leq k \leq n-1$, imply that $s \models C_n$.

Finally, in case s contains more than n worlds, pick an arbitrary element a and define the sets s^+, s^- as above. These sets have empty intersection and $s = s^+ \cup s^-$, thus $\#s = \#s^+ + \#s^-$. In particular, for every choice of a value $k \in \{1, \dots, n-1\}$ we have either that $\#s^+ > k$ or $\#s^- > n - k$. By the induction hypothesis, this implies that $s^+ \not\models C_k$ or $s^- \not\models C_{n-k}$. Since $s^+ \models_{[x \mapsto a]} P(x)$ and $s^- \models_{[x \mapsto a]} \neg P(x)$, it follows that $s \not\models_{[x \mapsto a]} P(x) \to C_k$ or $s \not\models_{[x \mapsto a]} \neg P(x) \to C_{n-k}$. Since the choices of a and $k \in \{1, \dots, n-1\}$ were arbitrary, this shows that $s \not\models C_n$.

To summarize, we showed that s supports the formula C_n if and only if s contains at most n worlds. This completes the inductive step. $\square$

Next, we show how the construction of the formulas C_n^Σ can be generalized to an arbitrary finite signature Σ which is function-rigid, i.e., such that all function symbols in Σ are rigid (this includes in particular all relational signatures, i.e., signatures without function symbols).

Definition 6.3 (Cardinality formulas, Σ function-rigid)**.**
Consider a finite function-rigid signature Σ. Let the predicate symbols in Σ be $R_1, \dots, R_l$. We define the cardinality formulas C_n^Σ inductively as follows, where

each $\overline{x}_i$ is a tuple of variables of size the arity of R_i.

$$\begin{array}{rcl} C_0^\Sigma & := & \bot \\ C_1^\Sigma & := & \forall \overline{x}_1 ?R_1(\overline{x}_1) \wedge \ldots \wedge \forall \overline{x}_l ?R_l(\overline{x}_l) \\ C_{n+1}^\Sigma & := & \exists\, \overline{x}_1 \bigvee\!\!\!\!\bigvee_{i=1}^{n} \big[(R_1(\overline{x}_1) \to C_i^\Sigma) \wedge (\neg R_1(\overline{x}_1) \to C_{n+1-i}^\Sigma) \big] \ \mathbin{\setminus\!\!\!/} \ldots \\ & \cdots \mathbin{\setminus\!\!\!/} & \exists\, \overline{x}_l \bigvee\!\!\!\!\bigvee_{i=1}^{n} \big[(R_l\,(\overline{x}_l) \to C_i^\Sigma) \wedge (\neg R_l\,(\overline{x}_l) \to C_{n+1-i}^\Sigma) \big] \end{array}$$

We can generalize this definition even further to allow for the possibility that Σ contains non-rigid function symbols, but in this case we must require that the identity predicate be available in Σ.

Definition 6.4 (Cardinality formulas, Σ including identity)**.**
Let Σ be a finite signature containing the identity predicate. Let $R_1, \ldots, R_l$ be the predicate symbols in Σ, and let $f_1, \ldots, f_h$ be the non-rigid function symbols in Σ. We define the cardinality formulas C_n^Σ inductively as follows, where $\overline{x}_j$ and $\overline{y}_j$ denote tuples of variables of size the arity of R_j and f_j respectively.

$$\begin{array}{rcl} C_0^\Sigma & := & \bot \\ C_1^\Sigma & := & \bigwedge_{j=1}^{l} \forall \overline{x}_j ?R_j(\overline{x}_j) \ \wedge \ \bigwedge_{j=1}^{h} \forall \overline{y}_j \exists z (f_j(\overline{y}_j) = z) \\ C_{n+1}^\Sigma & := & \bigvee\!\!\!\!\bigvee_{j=1}^{l} \exists \overline{x}_j \bigvee\!\!\!\!\bigvee_{i=1}^{n} [\, (R_j(\overline{x}_j) \to C_i^\Sigma) \wedge (\neg R_j(\overline{x}_j) \to C_{n+1-i}^\Sigma)\,] \ \mathbin{\setminus\!\!\!/} \\ & \mathbin{\setminus\!\!\!/} & \bigvee\!\!\!\!\bigvee_{j=1}^{h} \exists \overline{y}_j z \bigvee\!\!\!\!\bigvee_{i=1}^{n} [\, (f_j(\overline{y}_j) = z \ \to \ C_i^\Sigma) \ \wedge \ (f_j(\overline{y}_j) \neq z \ \to \ C_{n+1-i}^\Sigma)\,] \end{array}$$

The general version of Proposition 6.2 now reads as follows.

Proposition 6.5. Let Σ be a finite signature which is function-rigid or contains identity. Let M be a model for Σ, s a state in M and $n \in \mathbb{N}$. We have:

$$M, s \models C_n^\Sigma \iff \#_e s \leq n$$

The proof is a tedious but rather straightforward extension of the one of Proposition 6.2. We include it in Appendix A for the sake of completeness.

7 Bounded inquisitive logics

In this section, we consider a family of logics InqBQ_κ obtained by restricting the class of admissible models to those containing at most κ worlds. As we will see, studying the sequence of these logics allows us to draw some interesting conclusions about InqBQ.

First, we define an entailment relation $\models_{\mathsf{InqBQ}_\kappa}$ as follows:

$$\begin{array}{rcl} \Phi \models_{\mathsf{InqBQ}_\kappa} \psi & \iff & \text{for all models } M = \langle W, D, I \rangle \text{ with } \#W \leq \kappa, \\ & & \text{for all states } s \subseteq W \text{ and assignments } g : \\ & & M, s \models_g \Phi \text{ implies } M, s \models_g \psi \end{array}$$

We write InqBQ_κ for the set of formulas φ such that $\models_{\mathsf{InqBQ}_\kappa} \varphi$.

Note that if $\kappa < \lambda$ then the entailment relation $\models_{\mathsf{InqBQ}_\lambda}$ includes $\models_{\mathsf{InqBQ}_\kappa}$, and thus also $\mathsf{InqBQ}_\lambda \subseteq \mathsf{InqBQ}_\kappa$. So what we have is a monotonically shrinking sequence of logics, all of which include InqBQ and approximate the full logic more and more closely. A simple cardinality argument shows that InqBQ must coincide with one of these approximations.

Proposition 7.1. For every signature Σ, there is a cardinal κ such that $\models_{\mathsf{InqBQ}_\kappa}$ coincides with $\models_{\mathsf{InqBQ}}$.

Proof. We want to show that for some fixed cardinal κ, for every invalid entailment $\Phi \not\models \psi$, there exists a model M whose universe has cardinality at most κ, such that $M \models \Phi$ and $M \not\models \psi$. This suffices to prove the result.

Let $\{\langle \Phi_\alpha, \psi_\alpha \rangle \mid \alpha < \rho\}$ be an enumeration of the invalid entailments (i.e., $\Phi_\alpha \not\models \psi_\alpha$), where ρ is a suitably large cardinal. For every $\langle \Phi_\alpha, \psi_\alpha \rangle$ we can find an information model M_α such that $M_\alpha \models \varphi$ for every $\varphi \in \Phi_\alpha$ and $M_\alpha \not\models \psi_\alpha$. Define $\lambda_\alpha = \# W^{M_\alpha}$, the cardinality of the universe of M_α. Moreover, define $\lambda = \sup(\{\lambda_\alpha \mid \alpha < \rho\})$. All the models M_α have cardinality at most λ, thus all the invalid entailments of InqBQ are witnessed by some model of cardinality at most λ, as desired. $\square$

Another thing that we can say about our sequence of logics is that for finite cardinals, all inclusions are strict, i.e., we have

$$\mathsf{InqBQ}_0 \supsetneq \mathsf{InqBQ}_1 \supsetneq \mathsf{InqBQ}_2 \supsetneq \dots$$

To see this, it suffices to note that for all $n \in \mathbb{N}$ we have

$$C_n^{\{R\}} \in \mathsf{InqBQ}_n - \mathsf{InqBQ}_{n+1}$$

where R is an arbitrary predicate in Σ and $C_n^{\{R\}}$ the cardinality formula defined in the previous section.

In the setting of propositional inquisitive logic, an analogous sequence of approximations has been studied by Ciardelli (2009), who showed that propositional inquisitive logic InqB is the limit of its finitary approximations, in the sense that $\mathsf{InqB} = \bigcap_{n\in\mathbb{N}} \mathsf{InqB}_n$. A natural question, already raised by Sano (2011), is whether an analogous result holds in the first-order setting. The following proposition answers this question in the negative.

Proposition 7.2. $\mathsf{InqBQ} \neq \bigcap_{n\in\mathbb{N}} \mathsf{InqBQ}_n$

In words, this proposition says that InqBQ does not have the finite model property with respect to the universe: there exist sentences that are invalid, but that necessarily require an infinite universe in order to be refuted.

To prove this result, it is useful to first introduce a technical notion.

Definition 7.3. Given a predicate symbol $R \in \Sigma$, an information state s is an *R-chain* if for every pair of worlds $w, w' \in s$ we have $R_w \subseteq R_{w'}$ or $R_{w'} \subseteq R_w$.

	w_0	w_1	w_2	$\cdots$
a_0	■	■	■	
a_1	□	■	■	$\cdots$
a_2	□	□	■	
$\vdots$	$\vdots$	$\vdots$	$\vdots$	

Figure 3: An example of a P-chain, where we depict models using the same conventions as in Figure 2. The universe is the set $\{w_i \mid i \in \mathbb{N}\}$, while the domain is $\{a_i \mid i \in \mathbb{N}\}$. The extension of P is defined by the clause $a_i \in P_{w_j}$ iff $i \leq j$. To see that the model is a P-chain, observe that $P_{w_i} \subseteq P_{w_j}$ iff $i \leq j$.

In words, a state s is an R-chain if the extensions $\{R_w \mid w \in s\}$ are totally ordered by inclusion. An example of P-chain for a unary P is shown in Figure 3.

An important feature of R-chains is that they are definable in InqBQ.

Lemma 7.4. M is an R-chain $\iff$ $M \models \chi^R_{\text{chain}}$, where

$$\chi^R_{\text{chain}} \;:=\; \forall\overline{x}\,\forall\overline{y}\,[\,(R\overline{x} \to R\overline{y}) \vee\!\!\!\vee (R\overline{y} \to R\overline{x})\,]$$

Proof. First, suppose $M \not\models \chi^R_{\text{chain}}$. Let $\overline{a}, \overline{b}$ be tuples such that $M \not\models_{[\overline{x}\mapsto\overline{a},\overline{y}\mapsto\overline{b}]} R\overline{x} \to R\overline{y}$ and $M \not\models_{[\overline{x}\mapsto\overline{a},\overline{y}\mapsto\overline{b}]} R\overline{y} \to R\overline{x}$. Since both formulas are classical and thus truth-conditional (Proposition 2.3), there exist two worlds w, w' such that

$$M, w \models_{[\overline{x}\mapsto\overline{a}]} R\overline{x} \quad M, w \not\models_{[\overline{y}\mapsto\overline{b}]} R\overline{y} \qquad\qquad M, w' \not\models_{[\overline{x}\mapsto\overline{a}]} R\overline{x} \quad M, w' \models_{[\overline{y}\mapsto\overline{b}]} R\overline{y}$$

or equivalently

$$\overline{a} \in R_w \qquad \overline{b} \notin R_w \qquad\qquad \overline{a} \notin R_{w'} \qquad \overline{b} \in R_{w'} \tag{1}$$

Thus, $R_w \not\subseteq R_{w'}$ and $R_{w'} \not\subseteq R_w$, which shows that M is not an R-chain.

Conversely, suppose that M is not an R-chain. Then there exist two worlds w, w' such that $R_w \not\subseteq R_{w'}$ and $R_{w'} \not\subseteq R_w$; this means there exist two tuples $\overline{a}, \overline{b}$ for which (1) holds. From this it easily follows that $M \not\models \chi^R_{\text{chain}}$. □

We now have all the tools needed to prove Proposition 7.2.

Proof of Proposition 7.2. Take an arbitrary $R \in \Sigma$. We are going to show that the formula

$$\psi_R \;:=\; \chi^R_{\text{chain}} \to \exists\overline{x}(R\overline{x} \to \forall\overline{y}?R\overline{y})$$

is not valid in InqBQ but is valid in InqBQ_n for each n. For simplicity we spell out the proof for the case in which R is a unary predicate P, but the general case is analogous.

To see that $\psi_P \notin \mathsf{InqBQ}$, it suffices to verify that the model in Figure 3 is a counterexample. It remains to be shown that $\psi_P \in \mathsf{InqBQ}_n$ for every $n \in \mathbb{N}$, that is, that ψ_P is valid on all models with finitely many worlds.

Consider any finite state t such that $t \models \chi^P_{\text{chain}}$. By Lemma 7.4 we have that t is a P-chain, so the set $\{P_w \mid w \in t\}$ is totally ordered by $\subseteq$. Now there are two cases. If the set is a singleton, that means that the extension of P is the same at all worlds in t, so we have $t \models \forall y?Py$ and also $t \models \exists x(Px \to \forall y?Py)$. If the set $\{P_w \mid w \in t\}$ is not a singleton, then since it is finite (as t is finite) it must have a greatest element P^1 and a second greatest element P^2. Take any $a \in P^1 - P^2$, which ensures that P^1 is the only element in the set $\{P_w \mid w \in t\}$ containing a. Take a substate $t' \subseteq t$ such that $t' \models_{[x \mapsto a]} Px$. This means that at all $w \in t'$, the extension P_w contains a, and must thus coincide with P^1. Thus, the extension of P is the same at all $w \in t'$, which ensures $t' \models \forall y?Py$. Since t' was an arbitrary subset of t, this shows that $t \models \exists x(Px \to \forall y?Py)$.

We have thus shown that any finite state that supports the antecedent of ψ_P also supports the consequent, and thus that ψ_P is valid in all models with a finite universe. $\square$

The natural question to ask next is whether InqBQ coincides with its $\aleph_0$-approximation, obtained by restricting to models with countably many worlds. If so, this can be seen as a kind of downward Löwenheim-Skolem theorem with respect to the size fo the universe: every entailment which is invalid in InqBQ can be falsified in a model based on a countable universe. This is an important question that we will leave open here.

Open question. Does $\models_{\mathsf{InqBQ}}$ coincide with $\models_{\mathsf{InqBQ}_{\aleph_0}}$? If not, what is the least cardinal κ such that $\models_{\mathsf{InqBQ}}$ coincides with $\models_{\mathsf{InqBQ}_\kappa}$?

Note that the answers to these questions could in principle depend on the signature Σ.

8 Axiomatizing finite-bound inquisitive logics

In this section we are going to show that, given a finite signature Σ which is either function-rigid or contains identity, a strongly complete proof system for the logic InqBQ_n is obtained by extending the proof system for InqBQ shown in Figure 1 by means of the axiom C^Σ_n. More precisely, consider the proof system obtained by extending the system of Figure 1 with the following rule:

$$\frac{}{\vdash C^\Sigma_n}$$

We write $\Phi \vdash_n \psi$ if for some finite subset $\Phi_0 \subseteq \Phi$ the sequent $\Phi_0 \vdash \psi$ is derivable in this system. We write $\varphi \dashv\vdash_n \psi$ if φ and ψ are inter-derivable in this system, i.e., if $\varphi \vdash_n \psi$ and $\psi \vdash_n \varphi$. We are going to prove the following theorem.

Theorem 8.1 (Soundness and completeness for finite-bound inquisitive logics)**.** Let $\Phi \cup \{\psi\} \subseteq \mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ where Σ is a finite signature that is function-rigid or contains identity. Then:

$$\Phi \models_{\mathsf{InqBQ}_n} \psi \iff \Phi \vdash_n \psi$$

Soundness follows from the fact that each rule in Figure 1 is sound with respect to InqBQ_n, and C_n^Σ is valid in InqBQ_n. The rest of this section is devoted to showing completeness. We start from the simpler case in which Σ is a finite function-rigid signature, i.e., the case in which all function symbols in the signature are rigid. At the end of the section we discuss how to extend this to languages with non-rigid function symbols, in the presence of identity.

Before proceeding, we state two preliminary lemmas about derivability in the system $\vdash_n$. The straightforward proofs are left as exercises to the reader. We only remark that the $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$-split rule is needed for the third item of Lemma 8.3.

Lemma 8.2. For any formulas φ, ψ, χ, $(\varphi \leftrightarrow \psi) \vdash_n (\chi \leftrightarrow \chi')$, where χ' is any formula obtained from χ by replacing one or more occurrences of φ by ψ.

Lemma 8.3. For any formulas φ, ψ, χ and any classical formula α we have:

1. $\varphi \wedge (\psi \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} \chi) \dashv\vdash_n (\varphi \wedge \psi) \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} (\varphi \wedge \chi)$
2. $(\varphi \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} \psi) \to \chi \dashv\vdash_n (\varphi \to \chi) \wedge (\psi \to \chi)$
3. $\alpha \to (\varphi \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} \psi) \dashv\vdash_n (\alpha \to \varphi) \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} (\alpha \to \psi)$
4. $\exists x\varphi \to \psi \dashv\vdash_n \forall x(\varphi \to \psi)$, provided x is not free in ψ
5. $\varphi \to \forall x\psi \dashv\vdash_n \forall x(\varphi \to \psi)$, provided x is not free in φ

8.1 Saturated n-theories

Let Σ be a finite function-rigid signature, which will remain fixed throughout this section. To lighten the notation, we will write C_n for C_n^Σ. Let $\Sigma(A)$ be the extension of Σ with a countably infinite set A of fresh rigid constant symbols. We denote by Ter the set of rigid terms of the signature $\Sigma(A)$, by $\mathcal{L}^A$ the set of sentences in $\mathcal{L}_{\mathsf{InqBQ}}(\Sigma(A))$, by $\mathcal{L}_c^A$ the set of classical sentences in $\mathcal{L}^A$.

A crucial notion for our completeness proof is the notion of an *n-saturated theory.* An n-saturated theory is a set of sentences that has the right features to be the set of sentences supported by a non-empty information state of size at most n, in a model where the terms in Ter name all the individuals in the domain.

Definition 8.4 (Saturated n-theories)**.** A set of sentences $\Delta \subseteq \mathcal{L}^A$ is called a *saturated n-theory* over A if it has the following properties:

- Consistency: $\bot \notin \Delta$
- Deductive closure: for all $\varphi \in \mathcal{L}^A$, if $\Delta \vdash_n \varphi$ then $\varphi \in \Delta$
- Inquisitive disjunction property: if $\varphi \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} \psi \in \Delta$ then $\varphi \in \Delta$ or $\psi \in \Delta$

- Inquisitive existence property: if $\exists x\varphi \in \Delta$ then $\varphi[\mathsf{t}/x] \in \Delta$ for some $\mathsf{t} \in \mathsf{Ter}$
- Normality: if $\forall x\varphi \notin \Delta$ then $\varphi[\mathsf{t}/x] \notin \Delta$ for some $\mathsf{t} \in \mathsf{Ter}$

For convenience, we also introduce notions of derivability and equivalence relative to a saturated n-theory Δ, as follows:

- $\varphi \vdash_\Delta \psi \iff (\varphi \to \psi) \in \Delta$
- $\varphi \dashv\vdash_\Delta \psi \iff (\varphi \leftrightarrow \psi) \in \Delta$

The key property of saturated n-theories is given by the following lemma. The lemma says that given a saturated n-theory Δ, we can find formulas $\alpha_1, \dots \alpha_m$ with $m \leq n$ which, from the perspective of Δ, are exhaustive and mutually exclusive, and such that each α_i completely describes a single possible world, up to duplicates.

Lemma 8.5. If Δ is a saturated n-theory then for some $m \leq n$ there exist classical sentences $\alpha_1, \dots, \alpha_m \in \mathcal{L}_c^A$ such that:

1. $\alpha_1 \vee \cdots \vee \alpha_m \in \Delta$
2. $\neg(\alpha_i \wedge \alpha_j) \in \Delta$ for $i \neq j$
3. $\neg\alpha_i \notin \Delta$ for each $i \leq m$
4. $(\alpha_i \to C_1) \in \Delta$ for each $i \leq m$.

Proof. For simplicity we give the proof for the case $\Sigma = \{P^{(1)}\}$, but this extends straightforwardly to any function-rigid signature.

For $n = 1$ the statement is trivially satisfied by choosing $\alpha_1 = \top$, so we may assume that $n > 1$. Since Δ is a saturated n-theory and $\vdash_n C_n$, we have $C_n \in \Delta$, hence expanding the definition:

$$\begin{aligned}
&\exists x \textstyle\bigvee_{i=1}^{n-1}[(P(x) \to C_i) \wedge (\neg P(x) \to C_{n-i})] \in \Delta \\
&\qquad\quad \Downarrow \text{(by existence property, for some } \mathsf{t} \in \mathsf{Ter}) \\
&\textstyle\bigvee_{i=1}^{n-1}[(P(\mathsf{t}) \to C_i) \wedge (\neg P(\mathsf{t}) \to C_{n-i})] \in \Delta \\
&\qquad\quad \Downarrow \text{(by disjunction property, for some } k \text{ with } 1 \leq k < n) \\
&(P(\mathsf{t}) \to C_k) \wedge (\neg P(\mathsf{t}) \to C_{n-k}) \in \Delta \\
&\qquad\quad \Downarrow \text{(by deductive closure)} \\
&P(\mathsf{t}) \to C_k \in \Delta \quad \text{and} \quad \neg P(\mathsf{t}) \to C_{n-k} \in \Delta
\end{aligned}$$

We found that, for some $\mathsf{t} \in \mathsf{Ter}$ and some $k \in \{1, \dots, n-1\}$ we have $P(\mathsf{t}) \to C_k \in \Delta$ and $\neg P(\mathsf{t}) \to C_{n-k} \in \Delta$. Notice that the antecedents of the implications satisfy Conditions 1 and 2 of the statement, since we have $\vdash_n P(\mathsf{t}) \vee \neg P(\mathsf{t})$ and $\vdash_n \neg(P(\mathsf{t}) \wedge \neg P(\mathsf{t}))$ (recall that our proof system includes a complete proof system for classical first-order logic in restriction to classical formulas).

We can describe the situation in a more general way as follows: we have some classical formulas $\{\beta_1, \dots, \beta_l\}$ such that **(a)** $\vdash_n \beta_1 \vee \cdots \vee \beta_l$, **(b)** $\vdash_n \neg(\beta_i \wedge \beta_j)$

for $i \neq j$ and **(c)** there exist positive indices $h_1, \ldots, h_l$ such that $\sum_{j=1}^{l} h_j = n$ and $\beta_j \to C_{h_j} \in \Delta$ for every $j \leq l$.

Assuming we are given such a set $\{\beta_1, \ldots, \beta_l\}$, we are going to show how to replace an arbitrary formula β_j with associated index $h_j > 1$ with two new formulas β'_j and β''_j such that the resulting set $\{\beta_1, \ldots, \beta_{j-1}, \beta'_j, \beta''_j, \beta_{j+1}, \ldots, \beta_l\}$ still satisfies properties **(a)-(c)**.

Recall that $\beta_j \to C_{h_j} \in \Delta$ by property **(c)**. By expanding the definition of C_{h_j} in the expression we obtain:

$$
\begin{array}{l}
\beta_j \to \exists x \bigvee_{i=1}^{h_j-1}[(P(x) \to C_i) \wedge (\neg P(x) \to C_{h_j-i})] \in \Delta \\
\qquad\quad \Downarrow \text{(by } \exists\text{-split)} \\
\exists x \left[\; \beta_j \to \bigvee_{i=1}^{h_j-1}[(P(x) \to C_i) \wedge (\neg P(x) \to C_{h_j-i})] \; \right] \in \Delta \\
\qquad\quad \Downarrow \text{(by existence property, for some } \mathsf{t}' \in \mathsf{Ter}\text{)} \\
\beta_j \to \bigvee_{i=1}^{h_j-1}[(P(\mathsf{t}') \to C_i) \wedge (\neg P(\mathsf{t}') \to C_{h_j-i})] \in \Delta \\
\qquad\quad \Downarrow \text{(by } \bigvee\text{-split)} \\
\bigvee_{i=1}^{h_j-1}[\beta_j \to (P(\mathsf{t}') \to C_i) \wedge (\neg P(\mathsf{t}') \to C_{h_j-i})] \in \Delta \\
\qquad\quad \Downarrow \text{(by disjunction property, for some } k' \text{ with } 1 \leq k' < h_j\text{)} \\
\beta_j \to (P(\mathsf{t}') \to C_{k'}) \wedge (\neg P(\mathsf{t}') \to C_{h_j-k'}) \in \Delta \\
\qquad\quad \Downarrow \text{(by deductive closure)} \\
(\beta_j \wedge P(\mathsf{t}')) \to C_{k'} \in \Delta \quad \text{and} \quad (\beta_j \wedge \neg P(\mathsf{t}')) \to C_{h_j-k'} \in \Delta
\end{array}
$$

Thus, if we let $\beta'_j = (\beta_j \wedge P(\mathsf{t}'))$ and $\beta''_j = (\beta_j \wedge \neg P(\mathsf{t}'))$ we have that $\beta'_j \to C_{k'} \in \Delta$ and $\beta''_j \to C_{h_j-k'} \in \Delta$. Using this fact, it is easy to show that the resulting set $\{\beta_1, \ldots, \beta_{j-1}, \beta'_j, \beta''_j, \beta_{j+1}, \ldots, \beta_l\}$ still satisfies properties **(a)-(c)**.

We can then repeat this procedure until we end up with a set of n formulas $\{\alpha_1, \ldots, \alpha_n\}$ such that (a) $\vdash_n \alpha_1 \vee \cdots \vee \alpha_n$, (b) $\vdash_n \neg(\alpha_i \wedge \alpha_j)$ for $i \neq j$, and (c) $\alpha_i \to C_1 \in \Delta$ for $1 \leq i \leq n$. Thus, the set of formulas $\{\alpha_1, \ldots, \alpha_n\}$ obtained by this inductive procedure satisfies Conditions 1, 2 and 4 from the statement of the lemma. However, this set does not necessarily satisfy Condition 3, since we could have $\neg\alpha_i \in \Delta$ for some α_i in the set. To obtain the desired set, we simply remove the relevant α_i. Let us make this more precise.

Modulo reordering the formulas, we can assume that $\neg\alpha_1, \ldots, \neg\alpha_m \notin \Delta$ and that $\neg\alpha_{m+1}, \ldots, \neg\alpha_n \in \Delta$ for some $m \leq n$. We can now show that the formulas $\alpha_1, \ldots, \alpha_m$ satisfy all the required conditions:

1. For the first, condition, note that Δ contains $\neg\alpha_j$ for $j > m$ as well as the disjunction $\alpha_1 \vee \cdots \vee \alpha_n$. From these formulas, the disjunction $\alpha_1 \vee \cdots \vee \alpha_m$ is derivable. Since Δ is closed under deduction, $\alpha_1 \vee \cdots \vee \alpha_m \in \Delta$.
2. We know that $\vdash_n \neg(\alpha_i \wedge \alpha_j)$ for $i \neq j$, thus the second condition is satisfied.
3. The third condition is trivially satisfied, since the formulas $\alpha_1, \ldots, \alpha_m$ were chosen such that $\neg\alpha_1, \ldots, \neg\alpha_m \notin \Delta$.

4. By construction, Δ contains $\alpha_i \to C_1$ for every $i \leq n$, thus also for $i \leq m$.

This concludes our proof. $\square$

Given Δ, let us fix formulas $\alpha_1, \ldots, \alpha_m$ as in the previous lemma and call them the *world-sentences* for Δ. We denote the set of these world-sentences by $\mathcal{L}_w^\Delta$.

The fact that, from the perspecive of Δ, each world-sentence α_i completely describes a single possible world is brought out by the following lemma.

Lemma 8.6. Let Δ be a saturated n-theory and $\alpha_i \in \mathcal{L}_w^\Delta$. Then for any sentence $\varphi \in \mathcal{L}^A$ we have $(\alpha_i \to \varphi) \in \Delta$ or $(\alpha_i \to \neg\varphi) \in \Delta$.

Proof. The proof is by induction on φ. We only spell out the interesting cases, and as usual, we omit the superscript Σ to ease notation.

First suppose φ is an atomic sentence. Then φ is of the form $R\bar{\mathsf{t}}$, where $\bar{\mathsf{t}}$ is a sequence of rigid terms (this is because we are assuming for the moment that Σ, and thus also $\Sigma(A)$, is function-rigid). We know that Δ contains $\alpha_i \to C_1$. The sentence $\forall\overline{x}?R\overline{x}$ is a conjunct of C_1, so by deductive closure $(\alpha_i \to \forall\overline{x}?R\overline{x}) \in \Delta$. Since α_i is a sentence, $(\alpha_i \to \forall\overline{x}?R\overline{x})$ is provably equivalent to $\forall\overline{x}(\alpha_i \to ?R\overline{x})$, from which we can infer $\alpha_i \to ?R\bar{\mathsf{t}}$ since the terms in $\bar{\mathsf{t}}$ are all rigid. From this, since α_i is classical, by $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$-split we can derive $(\alpha_i \to R\bar{\mathsf{t}}) \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}} (\alpha_i \to \neg R\bar{\mathsf{t}})$. Thus, this formula is in Δ. Since Δ has the inquisitive disjunction property, it contains either $\alpha_i \to R\bar{\mathsf{t}}$ or $\alpha_i \to \neg R\bar{\mathsf{t}}$, as desired.

Suppose φ is an implication $\psi \to \chi$. By induction hypothesis, Δ contains either $\alpha_i \to \psi$ or $\alpha_i \to \neg\psi$, and moreover, it contains either $\alpha_i \to \chi$ or $\alpha_i \to \neg\chi$. If Δ contains either $\alpha_i \to \neg\psi$ or $\alpha_i \to \chi$, then it also contains $\alpha_i \to (\psi \to \chi)$ and we are done. Otherwise, Δ must contain both $\alpha_i \to \psi$ and $\alpha_i \to \neg\chi$, and then it also contains $\alpha_i \to \neg(\psi \to \chi)$.

Next, suppose φ is a universal formula $\forall x\psi$. By induction hypothesis, for each $\mathsf{t} \in \mathsf{Ter}$, Δ contains either $\alpha_i \to \psi[\mathsf{t}/x]$, or else $\alpha_i \to \neg\psi[\mathsf{t}/x]$. Now there are two cases: either Δ contains $\alpha_i \to \psi[\mathsf{t}/x]$ for all $\mathsf{t} \in \mathsf{Ter}$, or it contains $\alpha_i \to \neg\psi[\mathsf{t}/x]$ for some $\mathsf{t} \in \mathsf{Ter}$. In the former case, by normality Δ contains $\forall x(\alpha_i \to \psi)$ and thus, by deductive closure, also $\alpha_i \to \forall x\psi$. In the latter case, since for any t we have $\alpha_i \to \neg\psi[\mathsf{t}/x] \vdash_n \alpha_i \to \neg\forall x\psi$, it follows that Δ contains $\alpha_i \to \neg\forall x\psi$.

The case for the inquisitive existential quantifier is analogous, using the inquisitive existence property of Δ. $\square$

Next, we are going to show that from the perspective of a saturated n-theory Δ, each formula is equivalent to an inquisitive disjunction of classical formulas—more specifically, to an inquisitive disjunction of classical disjunctions of world-sentences for Δ.

To make this precise, we first extend the set of world-sentences to a set of state-sentences. These are sentences that, from the perspective of Δ, capture the possible information states.

Definition 8.7 (State sentences)**.** If $\mathcal{L}_w^\Delta = \{\alpha_1, \dots, \alpha_m\}$ is a set of world-sentences for a saturated n-theory Δ, the corresponding set of state-sentences is the set of classical disjunctions of world-sentences, with the addition of $\bot$:[8]

$$\mathcal{L}_s^\Delta = \{\bot\} \cup \{\alpha_{i_1} \vee \cdots \vee \alpha_{i_k} \mid 1 \le i_1 < \cdots < i_k \le m\}$$

Note that $\mathcal{L}_s^\Delta$ is finite.

The set of state sentences is closed under conjunction and implication, up to equivalence with respect to Δ. To illustrate the idea, suppose $\mathcal{L}_w^\Delta = \{\alpha_1, \alpha_2, \alpha_3\}$. Since the α_i are jointly exhaustive and mutually exclusive in Δ we have:

$$(\alpha_1 \vee \alpha_2) \wedge (\alpha_2 \vee \alpha_3) \dashv\vdash_\Delta \alpha_2 \qquad (\alpha_1 \vee \alpha_2) \to \bot \dashv\vdash_\Delta \alpha_3$$

The following lemma states the general fact. Since the proof is simply an exercise in classical logic, we leave it for Appendix B.

Lemma 8.8. For any two state sentences $\beta, \gamma \in \mathcal{L}_s^\Delta$ there are state-sentences $(\beta \sqcap \gamma)$ and $(\beta \sqsupset \gamma)$ in $\mathcal{L}_s^\Delta$ such that $(\beta \wedge \gamma) \dashv\vdash_\Delta (\beta \sqcap \gamma)$ and $(\beta \to \gamma) \dashv\vdash_\Delta (\beta \sqsupset \gamma)$.

We call $\beta \sqcap \gamma$ the *pseudo-conjunction* of β and γ, and $\beta \sqsupset \gamma$ the *pseudo-implication* of β and γ. If $S = \{\beta_1, \dots, \beta_k\}$ is a non-empty set of state-sentences, we also introduce the notation $\bigsqcap S$ for the state-sentence $\beta_1 \sqcap \cdots \sqcap \beta_k$.[9]

Next, we associate every sentence $\varphi \in \mathcal{L}^A$ with a finite set $\mathcal{R}_\Delta(\varphi) \subseteq \mathcal{L}_s^\Delta$ of state-sentences for Δ, which we call the Δ-*resolutions* of φ. The definition is largely parallel to the definition of resolutions for inquisitive propositional logic (see, e.g., Definition 2.4.1 in Ciardelli, 2016b), but it exploits the set of state-sentences for Δ in a crucial way.

Definition 8.9 (Δ-resolutions)**.** Let Δ be a saturated n-theory and $\mathcal{L}_w^\Delta$ a set of world-sentences for Δ. We define for each $\varphi \in \mathcal{L}^A$ a finite set $\mathcal{R}_\Delta(\varphi) \subseteq \mathcal{L}_s^\Delta$ as follows, where $\sqcap$ and $\sqsupset$ are the operations on state formulas given by Lemma 8.8:

- $\mathcal{R}_\Delta(p) = \{ \ \bigvee\{\alpha_i \in \mathcal{L}_w^\Delta \mid \alpha_i \vdash_\Delta p\} \ \}$, if p is an atomic sentence
- $\mathcal{R}_\Delta(\bot) = \{ \ \bot \ \}$
- $\mathcal{R}_\Delta(\varphi \wedge \psi) = \{ \ \beta \sqcap \gamma \mid \beta \in \mathcal{R}_\Delta(\varphi), \gamma \in \mathcal{R}_\Delta(\psi) \ \}$
- $\mathcal{R}_\Delta(\varphi \vee\!\!\!\vee \psi) = \mathcal{R}_\Delta(\varphi) \cup \mathcal{R}_\Delta(\psi)$
- $\mathcal{R}_\Delta(\varphi \to \psi) = \{ \ \bigsqcap_{\beta \in \mathcal{R}_\Delta(\varphi)} (\beta \sqsupset f(\beta)) \mid f : \mathcal{R}_\Delta(\varphi) \to \mathcal{R}_\Delta(\psi) \ \}$
- $\mathcal{R}_\Delta(\forall x \varphi) = \{ \ \bigsqcap\{f(\mathsf{t}) \mid \mathsf{t} \in \mathsf{Ter}\} \mid f \in \Pi_{\mathsf{t} \in \mathsf{Ter}} \mathcal{R}_\Delta(\varphi[\mathsf{t}/x]) \ \}$[10]
- $\mathcal{R}_\Delta(\exists x \varphi) = \bigcup_{\mathsf{t} \in \mathsf{Ter}} \mathcal{R}_\Delta(\varphi[\mathsf{t}/x])$

[8] Note that, if we stipulate that $\bigvee \emptyset = \bot$, can write $\mathcal{L}_s^\Delta = \{\bigvee S \mid S \subseteq \mathcal{L}_w^\Delta\}$.

[9] It can be checked based on the proof in the appendix that $\sqcap$ is associative and commutative, i.e., we have $(\beta \sqcap \gamma) \sqcap \delta = \beta \sqcap (\gamma \sqcap \delta)$ and $\beta \sqcap \gamma = \gamma \sqcap \beta$. Thus, the choice of bracketing in the definition of $\bigsqcap S$ and the order of the pseudo-conjuncts is not essential.

[10] Here, $\Pi_{\mathsf{t} \in \mathsf{Ter}_A} \mathcal{R}_\Delta(\varphi[\mathsf{t}/x])$ denotes the Cartesian product of the sets $\mathcal{R}_\Delta(\varphi[\mathsf{t}/x])$ for $\mathsf{t} \in \mathsf{Ter}$, that is, the set of functions f which associate to each $\mathsf{t} \in \mathsf{Ter}$ an element $f(\mathsf{t}) \in \mathcal{R}_\Delta(\varphi[\mathsf{t}/x])$.

It is immediate to check inductively that for each $\varphi \in \mathcal{L}^A$, $\mathcal{R}_\Delta(\varphi)$ is a subset of $\mathcal{L}_s^\Delta$, which is a finite set. Likewise, also the set $\{f(\mathsf{t}) \mid \mathsf{t} \in \mathsf{Ter}\}$ is a subset of $\mathcal{L}_s^\Delta$, and thus finite. This guarantees that the pseudo-conjunctions defined in the clause for implication and the universal quantifier are finitary and well-defined. With slight abuse of notation, we will also write the finite conjunction $\bigsqcap\{f(\mathsf{t}) \mid \mathsf{t} \in \mathsf{Ter}\}$ more compactly as $\bigsqcap_{\mathsf{t}\in\mathsf{Ter}} f(\mathsf{t})$, and we will similarly understand indexed conjunctions and disjunction as conjunctions and disjunctions of the corresponding (finite) set.

Next, we show that from the perspective of Δ, each formula is equivalent to the inquisitive disjunction of its Δ-resolutions. This is a version relativized to Δ of the normal form for inquisitive propositional logic (see Proposition 2.4.4 in Ciardelli, 2016b).

Lemma 8.10. For every $\varphi \in \mathcal{L}^A$, $\varphi \dashv\vdash_\Delta \mathbin{\rotatebox[origin=c]{180}{$\mathbin{\bigwedge\!\!\!\!\bigwedge}$}}$ $\mathcal{R}_\Delta(\varphi)$.

Proof. By induction on φ. We only give the most interesting inductive cases.

Suppose φ is an atom p. Then $\mathcal{R}_\Delta(p)$ is a singleton, and we need to show that $p \dashv\vdash_\Delta \alpha_{i_1} \vee \cdots \vee \alpha_{i_k}$ where $\alpha_{i_1} \ldots \alpha_{i_k}$ are all the world-sentences such that $\alpha_i \vdash_\Delta p$. In one direction, since each disjunct in $\alpha_{i_1} \vee \cdots \vee \alpha_{i_k}$ proves p on the basis of Δ, by classical reasoning we have $\alpha_{i_1} \vee \cdots \vee \alpha_{i_k} \vdash_\Delta p$ (note that all formulas involved are classical and recall that for classical formulas, our system includes a complete system for classical first-order logic). For the converse, take any α_j different from $\alpha_{i_1}, \ldots, \alpha_{i_k}$. By Lemma 8.6 we have $\alpha_j \vdash_\Delta \neg p$ and so by classical reasoning $p \vdash_\Delta \neg\alpha_j$. Recalling that $(\alpha_1 \vee \cdots \vee \alpha_n) \in \Delta$, it follows that $p \vdash_\Delta \alpha_{i_1} \vee \cdots \vee \alpha_{i_n}$.

Suppose φ is of the form $\psi \to \chi$. By inductive hypothesis $\psi \dashv\vdash_\Delta \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi)$ and $\chi \dashv\vdash_\Delta \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\chi)$. Thus by Lemma 8.2 we have:

$$\psi \to \chi \;\dashv\vdash_\Delta\; \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi) \;\to\; \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\chi)$$

By Lemmas 8.2 and 8.3 we have:

$$\begin{aligned}
\bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi) \;\to\; \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\chi) \quad &\dashv\vdash_\Delta \quad \textstyle\bigwedge_{\beta\in\mathcal{R}_\Delta(\psi)}(\beta \to \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\chi)) \\
&\dashv\vdash_\Delta \quad \textstyle\bigwedge_{\beta\in\mathcal{R}_\Delta(\psi)} \bigvee\!\!\!\bigvee_{\gamma\in\mathcal{R}_\Delta(\chi)} \beta \to \gamma \\
&\dashv\vdash_\Delta \quad \textstyle\bigvee\!\!\!\bigvee_{f:\mathcal{R}_\Delta(\psi)\to\mathcal{R}_\Delta(\chi)} \bigwedge_{\beta\in\mathcal{R}_\Delta(\psi)} \beta \to f(\beta)
\end{aligned}$$

For each $\beta \in \mathcal{R}_\Delta(\psi)$ and each function $f : \mathcal{R}_\Delta(\psi) \to \mathcal{R}_\Delta(\chi)$, the formulas β and $f(\beta)$ are in $\mathcal{L}_s^\Delta$; thus, by Lemmas 8.8 and 8.2 we have

$$\bigwedge_{\beta\in\mathcal{R}_\Delta(\psi)} \beta \to f(\beta) \dashv\vdash_\Delta \bigsqcap_{\beta\in\mathcal{R}_\Delta(\psi)} \beta \sqsupset f(\beta)$$

Putting together the previous equivalences, and using again Lemma 8.2, we have

$$\psi \to \chi \dashv\vdash_\Delta \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi \to \chi)$$

Next, suppose φ is of the form $\forall x\psi$. We first show that $\forall x\psi \vdash_\Delta \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi)$. By inductive hypothesis we have $\psi[\mathsf{t}/x] \vdash_\Delta \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x])$ for every $\mathsf{t} \in \mathsf{Ter}$, and thus also $\forall x\psi \vdash_\Delta \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x])$. Since this holds for every $\mathsf{t} \in \mathsf{Ter}$, we also have $\forall x\psi \vdash_\Delta \bigwedge_{\mathsf{t}\in\mathsf{Ter}} \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x])$. Note that the conjunction here is finite: since each $\mathcal{R}_\Delta(\psi[\mathsf{t}/x])$ is a set of state-sentences, and the state-sentences are finitely many, there are only finitely many such sets.

By the first item of Lemma 8.3 we have that

$$\bigwedge_{\mathsf{t}\in\mathsf{Ter}} \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x]) \dashv\vdash_\Delta \bigvee\!\!\!\!\bigvee \{ \bigwedge_{\mathsf{t}\in\mathsf{Ter}} f(\mathsf{t}) \mid f \in \Pi_{\mathsf{t}\in\mathsf{Ter}}\mathcal{R}_\Delta(\psi[\mathsf{t}/x])\}$$

For any function f and any $\mathsf{t} \in \mathsf{Ter}$, the formula $f(\mathsf{t})$ is a state-formula. Thus, by Lemma 8.8 we have $\bigwedge_{\mathsf{t}\in\mathsf{Ter}} f(\mathsf{t}) \dashv\vdash_\Delta \bigsqcap_{\mathsf{t}\in\mathsf{Ter}} f(\mathsf{t})$. Using Lemma 8.2, the right-hand side of the equation above is provably equivalent in Δ to

$$\bigvee\!\!\!\!\bigvee \{ \bigsqcap_{\mathsf{t}\in\mathsf{Ter}} f(\mathsf{t}) \mid f \in \Pi_{\mathsf{t}\in\mathsf{Ter}}\mathcal{R}_\Delta(\psi[\mathsf{t}/x])\}$$

which is nothing but $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi)$. This shows that $\forall x\psi \vdash_\Delta \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi)$.

We now prove the converse direction, $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi) \vdash_\Delta \forall x\psi$. Firstly notice that for every $\mathsf{t} \in \mathsf{Ter}$ and function $f \in \Pi_{\mathsf{t}'\in\mathsf{Ter}}\mathcal{R}_\Delta(\psi[\mathsf{t}'/x])$ we have $\bigsqcap_{\mathsf{t}'\in\mathsf{Ter}} f(\mathsf{t}') \vdash_\Delta f(\mathsf{t})$. Since $f(\mathsf{t})$ is an element of $\mathcal{R}_\Delta(\psi[\mathsf{t}/x])$, we have $\bigsqcap_{\mathsf{t}'\in\mathsf{Ter}} f(\mathsf{t}') \vdash_\Delta \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x])$. Since $\bigsqcap_{\mathsf{t}'\in\mathsf{Ter}} f(\mathsf{t}')$ is an arbitrary element of $\mathcal{R}_\Delta(\forall x\psi)$, we have

$$\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi) \vdash_\Delta \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x])$$

By inductive hypothesis $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x]) \vdash_\Delta \psi[\mathsf{t}/x]$, which combined with our previous conclusion leads us to $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi) \vdash_\Delta \psi[\mathsf{t}/x]$ for every $\mathsf{t} \in \mathsf{Ter}$. This means that for every $\mathsf{t} \in \mathsf{Ter}$, Δ contains the formula $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi) \to \psi[\mathsf{t}/x]$. By the normality condition, Δ must also contain $\forall x(\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi) \to \psi)$, which by item 5 of Lemma 8.3 is provably equivalent to $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi) \to \forall x\psi$. This means that $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi) \vdash_\Delta \forall x\psi$, concluding the inductive step for $\varphi = \forall x\psi$.

Finally, let φ be of the form $\exists x\psi$. We first prove that $\exists x\psi \vdash_\Delta \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi)$. By inductive hypothesis we have that $\psi[\mathsf{t}/x] \vdash_\Delta \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x])$ for every $\mathsf{t} \in \mathsf{Ter}$. Since $\mathcal{R}_\Delta(\psi[\mathsf{t}/x]) \subseteq \mathcal{R}_\Delta(\exists x\psi)$, we have $\psi[\mathsf{t}/x] \vdash_\Delta \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi)$. This means that Δ contains the formula $\psi[\mathsf{t}/x] \to \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi)$ for every $\mathsf{t} \in \mathsf{Ter}$, and so by normality it must also contain $\forall x(\psi \to \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi))$. By item 4 of Lemma 8.3, this is provably equivalent to $\exists x\psi \to \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi)$, which must then be in Δ. This means that $\exists x\psi \vdash_\Delta \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi)$.

We now prove that $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi) \vdash_\Delta \exists x\psi$. By inductive hypothesis we have that for every $\mathsf{t} \in \mathsf{Ter}$, $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x]) \vdash_\Delta \psi[\mathsf{t}/x]$, and so also $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x]) \vdash_\Delta \exists x\psi$ (note that $\psi[\mathsf{t}/x] \vdash_n \exists x\psi$ since t is rigid). Since this is the case for each $\mathsf{t} \in \mathsf{Ter}$ we also have $\bigvee\!\!\!\!\bigvee_{\mathsf{t}\in\mathsf{Ter}} \bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x]) \vdash_\Delta \exists x\psi$. Finally, since $\mathcal{R}_\Delta(\exists x\psi) = \bigcup_{\mathsf{t}\in\mathsf{Ter}} \mathcal{R}_\Delta(\psi[\mathsf{t}/x])$, the left-hand side coincides with $\bigvee\!\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi)$. This concludes the inductive step for $\varphi = \exists x\psi$. $\square$

This shows that within the context of a saturated n-theory Δ, every sentence is provably equivalent to one of a finite set of sentences which have a special form: they are inquisitive disjunctions of classical disjunctions of world-sentences.

8.2 Canonical model

Building on the results in the previous section, we will now show that for any saturated n-theory Δ there is a canonical model M^c_Δ, based on a universe of at most n worlds, that supports all and only the sentences in Δ. We can define the canonical model for Δ on the basis of the world-sentences for Δ, as follows.

Definition 8.11 (Canonical model)**.**
Let Δ be a saturated n-theory and $\mathcal{L}^\Delta_w = \{\alpha_1, \dots, \alpha_m\}$ a set of world-sentences for Δ. The canonical model for Δ is $M^c_\Delta = \langle W^c, D^c, I^c \rangle$ where:

- $W^c = \{w_1, \dots, w_m\}$ is an arbitrary set of m worlds (where $m \leq n$ is the number of world-sentences for Δ);
- D^c is Ter, the set of closed rigid terms in the signature $\Sigma(A)$;[11]
- I^c is defined as follows:
  - if f is a rigid function symbol, $I^c(\mathsf{f})(\mathsf{t}_1, \dots, \mathsf{t}_l)$ is the term $\mathsf{f}(\mathsf{t}_1, \dots, \mathsf{t}_l)$; in particular, if c is a rigid constant then $I^c(\mathsf{c}) = \mathsf{c}$;[12]
  - if R is a predicate symbol and $w_i \in W^c$,

$$\langle \mathsf{t}_1, \dots, \mathsf{t}_l \rangle \in I^c_{w_i}(R) \iff \alpha_i \vdash_\Delta R(\mathsf{t}_1, \dots, \mathsf{t}_l)$$

A straightforward induction shows that in this model, every rigid closed term $\mathsf{t} \in \mathsf{Ter}$ is interpreted rigidly as itself.

Lemma 8.12. For each $\mathsf{t} \in \mathsf{Ter}$, world $w_i \in W^c$ and assignment g, $[\mathsf{t}]^g_{w_i} = \mathsf{t}$.

Using this fact, it is routine to show that quantification can be handled by substitution, in the following sense.

Lemma 8.13. Let s be any state of M^c_Δ and let φ be a formula where at most x occurs free. We have:

$$\begin{array}{lcl} M^c_\Delta, s \models \forall x \varphi & \iff & M^c_\Delta, s \models \varphi[\mathsf{t}/x] \text{ for all } \mathsf{t} \in D^c \\ M^c_\Delta, s \models \exists x \varphi & \iff & M^c_\Delta, s \models \varphi[\mathsf{t}/x] \text{ for some } \mathsf{t} \in D^c \end{array}$$

Next we show a truth lemma restricted to classical sentences.

Lemma 8.14 (Truth lemma for classical sentences)**.**
For every classical sentence $\gamma \in \mathcal{L}^A_c$ and every world $w_i \in W^c$ we have:

$$M^c_\Delta, w_i \models \gamma \iff \alpha_i \vdash_\Delta \gamma$$

Proof. The proof is by induction on γ. The atomic case is immediate by Lemma 8.12 and the interpretation of predicate symbols. The case for $\bot$ follows since world-sentences are chosen such that $\neg\alpha_i \notin \Delta$, i.e., such that $\alpha_i \not\vdash_\Delta \bot$. The inductive step for conjunction is immediate. We spell out the remaining two cases, for $\to$ and $\forall$.

[11] Here, this coincides with the set of all closed terms, since for now our signature is assumed to be function-rigid, and thus all terms are rigid. However, the definition is formulated with an eye to the extension to non-rigid function symbols, discussed in Subsection 8.4.

[12] This is the only case that we need to consider for now, since we are assuming that all function symbols in our signature are rigid. We will relax this assumption in Subsection 8.4.

- Suppose $\gamma = \delta \to \eta$. We have:

$$\begin{aligned} M^c_\Delta, w_i \models \delta \to \eta \quad &\iff \quad M^c_\Delta, w_i \not\models \delta \text{ or } M^c_\Delta, w_i \models \eta \\ &\iff \quad \alpha_i \nvdash_\Delta \delta \text{ or } \alpha_i \vdash_\Delta \eta \\ &\iff \quad \alpha_i \vdash_\Delta \neg\delta \text{ or } \alpha_i \vdash_\Delta \eta \\ &\iff \quad \alpha_i \vdash_\Delta (\delta \to \eta) \end{aligned}$$

  Here, the second equivalence is the induction hypothesis. The third equivalence is given by Lemma 8.6. As for the last equivalence, the left-to-right direction is clear. For the converse, suppose $\alpha_i \vdash_\Delta \delta \to \eta$ and suppose towards a contradiction that $\alpha_i \nvdash_\Delta \neg\delta$ and $\alpha_i \nvdash_\Delta \eta$. Then by Lemma 8.6 we have $\alpha_i \vdash_\Delta \delta$ and $\alpha_i \vdash_\Delta \neg\eta$. Now from $\alpha_i \vdash_\Delta \delta \to \eta$ and $\alpha_i \vdash_\Delta \delta$ we have $\alpha_i \vdash_\Delta \eta$, which together with $\alpha_i \vdash_\Delta \neg\eta$ implies $\alpha_i \vdash_\Delta \bot$, i.e., $\neg\alpha_i \in \Delta$. This is a contradiction since the world-formulas are chosen in such a way that $\neg\alpha_i \notin \Delta$.

- Suppose $\gamma = \forall x\delta$. We have:

$$\begin{aligned} M^c_\Delta, w_i \models \forall x\delta \quad &\iff \quad M^c_\Delta, w_i \models \delta[\mathsf{t}/x] \ \text{ for all } \mathsf{t} \in D^c \\ &\iff \quad \alpha_i \vdash_\Delta \delta[\mathsf{t}/x] \ \text{ for all } \mathsf{t} \in D^c \\ &\iff \quad \alpha_i \vdash_\Delta \forall x\delta \end{aligned}$$

  Here, the first equivalence is given by Lemma 8.13. The second equivalence is the induction hypothesis. The right-to-left direction of the last equivalence is given by the fact the terms $\mathsf{t} \in D^c$ are rigid, and so $\forall x\delta \vdash_\Delta \delta[\mathsf{t}/x]$. For the converse, recall that $\alpha_i \vdash_\Delta \delta[\mathsf{t}/x]$ means that $(\alpha_i \to \delta[\mathsf{t}/x]) \in \Delta$. By the normality condition, if this is the case for all closed terms then also $\forall x(\alpha_i \to \delta) \in \Delta$. Since α_i is a sentence, by Lemma 8.3 this is interderivable with $\alpha_i \to \forall x\delta$, and so the latter formula is in Δ, which means that $\alpha_i \vdash_\Delta \forall x\delta$.

□

The next step on the way to our conclusion is to prove that the normal form given by Lemma 8.10 is semantically sound in the canonical model. We start with the following lemma.

Lemma 8.15. Consider two state-sentences $\beta, \gamma \in \mathcal{L}^\Delta_s$ and the operators $\sqcap, \sqsupset$ introduced in Lemma 8.8. We have that $M^c_\Delta \models (\beta \wedge \gamma) \leftrightarrow (\beta \sqcap \gamma)$ and $M^c_\Delta \models (\beta \to \gamma) \leftrightarrow (\beta \sqsupset \gamma)$.

Proof. We show that $M^c_\Delta \models (\beta \wedge \gamma) \leftrightarrow (\beta \sqcap \gamma)$ (the proof of the second claim is analogous). The formula $(\beta \wedge \gamma) \leftrightarrow (\beta \sqcap \gamma)$ is classical and thus truth-conditional, so it suffices to show that $M^c_\Delta, w_i \models (\beta \wedge \gamma) \leftrightarrow (\beta \sqcap \gamma)$ for every world $w_i \in W^c$. By Lemma 8.14 this is equivalent to showing that $\alpha_i \vdash_\Delta (\beta \wedge \gamma) \leftrightarrow (\beta \sqcap \gamma)$ for every world-sentence α_i. The last statement follows immediately from Lemma 8.8. □

We are now ready to show soundness of the normal form.

Lemma 8.16. For every $\varphi \in \mathcal{L}^A$ and any state s of the canonical model M^c_Δ we have $M^c_\Delta, s \models \varphi \iff M^c_\Delta, s \models \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\varphi)$.

Proof. By induction on φ. We only give the most interesting inductive cases, namely, the ones for atoms, implication, and the quantifiers. Throughout the proof we indicate with IH the inductive hypothesis and drop reference to M^c_Δ.

- φ is an atomic sentence p. In this case $\mathcal{R}_\Delta(p)$ contains only one element, $p' := \bigvee\{\alpha_i \in \mathcal{L}^\Delta_w \mid \alpha_i \vdash_\Delta p\}$ and $\bigvee\!\!\!\bigvee \mathcal{R}_\Delta(p) = p'$. By Lemma 8.10 we have $p \dashv\vdash_\Delta p'$. Thus, for any world-sentence α_i we have $\alpha_i \vdash_\Delta p \iff \alpha_i \vdash_\Delta p'$. Since p and p' are classical, the truth-lemma implies that p and p' are true at the same worlds in M^c_Δ, and thus supported at the same states.
- $\varphi = (\psi \to \chi)$. For the left-to-right direction, take any state s in M^c_Δ and suppose $s \models \psi \to \chi$. Take an arbitrary $\alpha \in \mathcal{R}_\Delta(\psi)$ and let $t_\alpha := \{w \in s \mid w \models \alpha\}$. Since α is classical and thus truth-conditional, we have $t_\alpha \models \alpha$. By IH, this implies that $t_\alpha \models \psi$. Since $s \models \psi \to \chi$, it follows that $t_\alpha \models \chi$. Again by IH, we get $t_\alpha \models \beta$ for some resolution $\beta \in \mathcal{R}_\Delta(\chi)$. Since α is truth-conditional, any $t \subseteq s$ that supports α is included in t_α, and thus by persistency supports β. This ensures that $s \models \alpha \to \beta$. By Lemma 8.15, it follows that $s \models \alpha \sqsupset \beta$. We have thus shown that for every $\alpha \in \mathcal{R}_\Delta(\psi)$ there is a $\beta \in \mathcal{R}_\Delta(\chi)$ such that $s \models \alpha \sqsupset \beta$. Hence, there is a function $f : \mathcal{R}_\Delta(\psi) \to \mathcal{R}_\Delta(\chi)$ such that $s \models \bigwedge_{\alpha \in \mathcal{R}_\Delta(\psi)}(\alpha \to f(\alpha))$. By Lemma 8.15, this is equivalent to $s \models \bigsqcap_{\alpha \in \mathcal{R}_\Delta(\psi)}(\alpha \to f(\alpha))$. By definition, this pseudo-conjunction is an element of $\mathcal{R}_\Delta(\psi \to \chi)$, so $s \models \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi \to \chi)$.

  For the converse, suppose that $s \models \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi \to \chi)$. Then there is some $f : \mathcal{R}_\Delta(\psi) \to \mathcal{R}_\Delta(\chi)$ such that $s \models \bigsqcap_{\alpha \in \mathcal{R}_\Delta(\psi)}(\alpha \sqsupset f(\alpha))$, which by Lemma 8.15 is equivalent to $s \models \bigwedge_{\alpha \in \mathcal{R}_\Delta(\psi)}(\alpha \to f(\alpha))$. We want to show that $s \models \psi \to \chi$. So, take any $t \subseteq s$ and suppose $t \models \psi$. By IH, $t \models \alpha$ for some $\alpha \in \mathcal{R}_\Delta(\psi)$. Since $s \models \alpha \to f(\alpha)$ and $t \subseteq s$ we have $t \models f(\alpha)$. Since $f(\alpha) \in \mathcal{R}_\Delta(\chi)$, by IH we have $t \models \chi$. This shows that $s \models \psi \to \chi$, as desired.
- $\varphi = \forall x \psi$. We have:

$$
\begin{array}{lll}
 & M^c_\Delta, s \models \forall x\psi & \\
\iff & \text{for all } \mathsf{t} \in \mathsf{Ter},\ M^c_\Delta, s \models \psi[\mathsf{t}/x] & \text{(by normality)} \\
\iff & \text{for all } \mathsf{t} \in \mathsf{Ter},\ M^c_\Delta, s \models \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x]) & \text{(by IH)} \\
\iff & s \models \bigwedge\{\ \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x]) \mid \mathsf{t} \in \mathsf{Ter}\ \} & \\
\iff & s \models \bigvee\!\!\!\bigvee\{\ \bigwedge f(\mathsf{t}) \mid f \in \Pi_{\mathsf{t}\in\mathsf{Ter}} \mathcal{R}_\Delta(\psi[\mathsf{t}/x])\ \} & \\
\iff & s \models \bigvee\!\!\!\bigvee\{\ \bigsqcap f(\mathsf{t}) \mid f \in \Pi_{\mathsf{t}\in\mathsf{Ter}} \mathcal{R}_\Delta(\psi[\mathsf{t}/x])\ \} & \text{(by Lemma 8.15)} \\
\iff & s \models \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\forall x\psi) & \text{(by def. of } \mathcal{R}_\Delta(\forall x\psi))
\end{array}
$$

- $\varphi = \exists x \psi$. We have:

$$
\begin{array}{lll}
 & M^c_\Delta, s \models \exists x\psi & \\
\iff & \text{for some } \mathsf{t} \in \mathsf{Ter},\ M^c_\Delta, s \models \psi[\mathsf{t}/x] & \\
\iff & \text{for some } \mathsf{t} \in \mathsf{Ter},\ M^c_\Delta, s \models \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\psi[\mathsf{t}/x]) & \text{(by IH)} \\
\iff & s \models \bigvee\!\!\!\bigvee \mathcal{R}_\Delta(\exists x\psi) & \text{(by def. of } \mathcal{R}_\Delta(\exists x\psi))
\end{array}
$$

□

It will also be handy to have remarked the following fact explicitly.

Lemma 8.17. For all $\beta \in \mathcal{L}_c^A$, $\beta \in \Delta \iff (\alpha_i \vdash_\Delta \beta$ for each $i = 1, \ldots, m)$.

Proof. Recall from Lemma 8.5 that $(\alpha_1 \vee \cdots \vee \alpha_m) \in \Delta$, and so using classical reasoning we have:

$$\begin{aligned} \beta \in \Delta \quad &\iff \quad \alpha_1 \vee \cdots \vee \alpha_m \to \beta \in \Delta \\ &\iff \quad (\alpha_1 \to \beta) \wedge \cdots \wedge (\alpha_n \to \beta) \in \Delta \\ &\iff \quad (\alpha_1 \to \beta) \in \Delta \text{ and } \ldots \text{ and } (\alpha_n \to \beta) \in \Delta \end{aligned}$$

□

Finally, with these results at hand, we are ready to show that the set of sentences supported by the canonical model M_Δ^c coincides precisely with Δ.

Lemma 8.18 (Support lemma)**.**
For every sentence $\varphi \in \mathcal{L}^A$ we have $M_\Delta^c \models \varphi \iff \varphi \in \Delta$.

Proof. By using the results collected in the current section, we obtain:

$$\begin{aligned} M_\Delta^c \models \varphi \quad &\iff \quad M_\Delta^c \models \bigvee\!\!\!\!\!\bigvee \mathcal{R}_\Delta(\varphi) && \text{(by Lemma 8.16)} \\ &\iff \quad \exists \beta \in \mathcal{R}_\Delta(\varphi) : M_\Delta^c \models \beta && \\ &\iff \quad \exists \beta \in \mathcal{R}_\Delta(\varphi)\ \forall w_i \in W^c : M_\Delta^c, w_i \models \beta && \\ &\iff \quad \exists \beta \in \mathcal{R}_\Delta(\varphi)\ \forall i \leq m : \alpha_i \vdash_\Delta \beta && \text{(by Lemma 8.14)} \\ &\iff \quad \exists \beta \in \mathcal{R}_\Delta(\varphi) : \beta \in \Delta && \text{(by Lemma 8.17)} \\ &\iff \quad \bigvee\!\!\!\!\!\bigvee \mathcal{R}_\Delta(\varphi) \in \Delta && \text{(by inq. disj. prop.)} \\ &\iff \quad \varphi \in \Delta && \text{(by Lemma 8.10)} \end{aligned}$$

□

This shows that every saturated n-theory coincides with the set of sentences which are supported by a model with at most n worlds.

8.3 Completeness

Let us now see how our canonical model construction can be used to establish the completeness of our system $\vdash_n$. First, we need a saturation lemma. Since the proof of this result uses a standard saturation argument (cf. Gabbay, 1981), we leave the details for a technical appendix (see Appendix C).

Lemma 8.19 (Saturation lemma)**.**
Suppose $\Phi \cup \{\psi\}$ is a set of sentences in the signature Σ with $\Phi \nvdash_n \psi$. Then there is a saturated n-theory Δ in the language $\mathcal{L}^A$ such that $\Phi \subseteq \Delta$ and $\psi \notin \Delta$.

With this saturation lemma in place, we are now ready to prove completeness. We first do so for the case of sentences.

Theorem 8.20. Suppose $\Phi \cup \{\psi\}$ be a set of sentences in the function-rigid signature Σ. If $\Phi \models_{\mathsf{InqBQ}_n} \psi$, then $\Phi \vdash_n \psi$.

Proof. By contraposition, suppose $\Phi \nvdash_n \psi$. By the saturation lemma we can find a saturated n-theory Δ in the extended language $\mathcal{L}^A$ with $\Phi \subseteq \Delta$ and $\psi \notin \Delta$. By Lemma 8.18 there is a model M based on a universe of at most n worlds such that Δ is exactly the set of sentences supported by M. In particular, M supports all formulas in Φ but not ψ, which shows that $\Phi \not\models_{\mathsf{InqBQ}_n} \psi$. □

Finally, we can extend this result to open formulas, thus proving Theorem 8.1 for the case of a function-rigid signature.

Proof of Theorem 8.1, case for function-rigid signatures. Suppose $\Phi \cup \{\psi\}$ is a set of formulas in a finite function-rigid signature Σ such that $\Phi \nvdash_n \psi$. Let Σ^* be a larger signature obtained by adding a rigid constant c_x for each variable x occurring free in $\Phi \cup \{\psi\}$. Let $\Phi^* \cup \{\psi^*\}$ be the set of sentences obtained from $\Phi \cup \{\psi\}$ by replacing each free occurrence of x by c_x. We have that $\Phi^* \nvdash_n \psi^*$ (if we had $\Phi^* \vdash_n \psi^*$, by a simple substitution we could turn a proof of this into a proof of $\Phi \vdash_n \psi$). Thus by the proof of the previous theorem, we have a model M based on a universe of at most n worlds and on a domain of closed terms including all constants c_x such that $M \models \Phi^*$ but $M \not\models \psi^*$. Then defining an assignment g such that $g(x) = \mathsf{c}_x$ for all variables x we have $M \models_g \Phi$ and $M \not\models_g \psi$. Since M has at most n worlds, this shows that $\Phi \not\models_{\mathsf{InqBQ}_n} \psi$. □

8.4 Adding non-rigid function symbols

Throughout this section, we have so far assumed all function symbols to be rigid. We now show how to deal with the slight complications that arise if the signature Σ contains non-rigid function symbols, provided identity is available.

The key to the generalization lies in the following lemma, that says that a saturated theory identifies each closed term with a rigid term.

Lemma 8.21. Let Δ be a saturated n-theory and $\alpha_i \in \mathcal{L}_w^\Delta$ a world-sentence for Δ. Then for any closed term t in the signature $\Sigma(A)$, we have $\alpha_i \vdash_\Delta (t = \mathsf{t})$ for some rigid term $\mathsf{t} \in \mathsf{Ter}$.

Proof. By induction over the structure of t. We only spell out the most interesting case, namely, the inductive step for t of the form $f(t_1, \ldots, t_k)$ for some non-rigid function symbol f.

By induction hypothesis, for each t_j we have a corresponding rigid term t_j such that $\alpha_i \vdash_\Delta (t_j = \mathsf{t}_j)$. Since f is non-rigid, the formula $\forall \overline{x} \exists y (f(\overline{x}) = y)$ is a conjunct of C_1^Σ. Since $\alpha_i \vdash_\Delta C_1^\Sigma$, by deductive closure of Δ we have also $\alpha_i \vdash_\Delta \forall \overline{x} \exists y (f(\overline{x}) = y)$. This means that Δ contains the formula $\alpha_i \to \forall \overline{x} \exists y (f(\overline{x}) = y)$, and since α_i is a sentence, this is equivalent to $\forall \overline{x} (\alpha_i \to \exists y (f(\overline{x}) = y))$. Instantiating the universal quantifier with $\mathsf{t}_1, \ldots, \mathsf{t}_k$, which is possible since these terms are rigid, we obtain $(\alpha_i \to \exists y (f(\mathsf{t}_1, \ldots, \mathsf{t}_k) = y)) \in \Delta$. By the $\exists$-split rule it follows that $\exists y (\alpha_i \to (f(\mathsf{t}_1, \ldots, \mathsf{t}_k) = y)) \in \Delta$, and thus by the inquisitive existence property that $\alpha_i \to (f(\mathsf{t}_1, \ldots, \mathsf{t}_k) = \mathsf{t}) \in \Delta$ for some

$\mathsf{t} \in \mathsf{Ter}$. Finally by the rules for identity we obtain $\alpha_i \to (f(t_1, \ldots, t_k) = \mathsf{t}) \in \Delta$, which means that $\alpha_i \vdash_\Delta (f(t_1, \ldots, t_k) = \mathsf{t})$ for some rigid t. □

Notice that this result relies crucially on the presence of the identity predicate in the signature. For this reason, in the presence of non-rigid symbols we require identity to be available.

We now need to supplement some of the definitions and proofs presented in this section. We start with the proof of Lemma 8.6.

Proof of Lemma 8.6, addendum. We now need to consider the case of atoms of the form $R\bar{t}$ for $\bar{t}$ a sequence of closed terms, which are possibly non-rigid. We need to show that, for each world-sentence $\alpha_i \in \mathcal{L}_w^\Delta$ we have $(\alpha_i \to R\bar{t}) \in \Delta$ or $(\alpha_i \to \neg R\bar{t}) \in \Delta$.

Let $\bar{t} = \langle t_1, \ldots, t_l \rangle$. By Lemma 8.21, we have a sequence $\bar{\mathsf{t}} = \langle \mathsf{t}_1, \ldots, \mathsf{t}_l \rangle$ of rigid terms in Ter such that $\alpha_i \vdash_\Delta t_j = \mathsf{t}_j$ for each $j \leq l$. By the rules for identity we can infer that $\alpha_i \to (R\bar{t} \leftrightarrow R\bar{\mathsf{t}}) \in \Delta$. Notice that $\mathsf{t}_1, \ldots, \mathsf{t}_l$ are rigid terms in the signature $\Sigma(A)$, so we already showed in the original proof of Lemma 8.6 that we have $(\alpha_i \to R\bar{\mathsf{t}}) \in \Delta$ or $(\alpha_i \to \neg R\bar{\mathsf{t}}) \in \Delta$. Thus combining these two facts we can conclude that $(\alpha_i \to R\bar{t}) \in \Delta$ or $(\alpha_i \to \neg R\bar{t}) \in \Delta$. □

Next, we need to supplement the definition of the canonical model M_Δ^c (Definition 8.11) with the interpretation of non-rigid function symbols. Whereas the interpretation of a rigid function symbol is uniform throughout the model, the interpretation of a non-rigid function symbol may vary from world to world.

Definition 8.22 (Addendum to Definition 8.11)**.** In defining the canonical model, we augment the definition of the interpretation I^c with the following clause:

- if f is a non-rigid function symbol and $w_i \in W^c$, $I_{w_i}^c(f)(\mathsf{t}_1, \ldots, \mathsf{t}_l) = \mathsf{t}$ for some term $\mathsf{t} \in \mathsf{Ter}$ such that $\alpha_i \vdash_\Delta f(\mathsf{t}_1, \ldots, \mathsf{t}_l) = \mathsf{t}$.

Notice that the existence of such a term t is guaranteed by Lemma 8.21. In case several $\mathsf{t} \in \mathsf{Ter}$ satisfy the property, we may choose arbitrarily among them.

A straightforward inductive proof now yields the following generalization of Lemma 8.12.

Lemma 8.23. For each closed term t, world $w_i \in W^c$ and assignment g, $[t]_{w_i}^g = \mathsf{t}$ for some $\mathsf{t} \in \mathsf{Ter}$ such that $\alpha_i \vdash_\Delta (t = \mathsf{t})$.

Finally, we need to augment the proof of the truth lemma to cover atomic formulas involving non-rigid terms.

Proof of Lemma 8.14, addendum. We now need to consider atoms of the form $R(t_1, \ldots, t_l)$ where $t_1, \ldots, t_l$ are arbitrary closed terms. Let $\mathsf{t}_j = [t_j]_{w_i}$. By Lemma 8.23 we know that $\alpha_i \vdash_\Delta (t_j = \mathsf{t}_j)$. By the interpretation of predicates in the canonical model and by the inference rules for identity we have

$$\begin{aligned} M^c_\Delta, w_i \models R(t_1, \dots, t_l) &\iff \langle [t_1]_{w_i}, \dots, [t_l]_{w_i} \rangle \in R_{w_i} \\ &\iff \langle \mathsf{t}_1, \dots, \mathsf{t}_l \rangle \in R_{w_i} \\ &\iff \alpha_i \vdash_\Delta R(\mathsf{t}_1, \dots, \mathsf{t}_l) \\ &\iff \alpha_i \vdash_\Delta R(t_1, \dots, t_l) \end{aligned}$$

as desired. □

9 A complete proof system for entailments with rex conclusions

The results in the previous section can be used to provide a proof system for InqBQ which can be shown to be complete for entailments with a rex conclusion, provided the signature Σ contains identity or is function-rigid (we do not, however, require Σ to be finite). Given such a signature Σ, we extend the natural deduction system in Figure 1 with the following inference rule, which we will refer to as the *coherence rule*:

$$\frac{\Theta, C_n^{\Sigma'} \vdash \chi}{\Theta \vdash \chi}$$

where we require $\chi \in \mathcal{L}_{\mathsf{rex}}$ and where $C_n^{\Sigma'}$ is any cardinality formula obtained for some number $n \geq n_\chi$ and some signature $\Sigma' \subseteq \Sigma$ (thus, Σ' must be finite and must be function-rigid or contain identity). In words, the rule allows us to discharge a cardinality assumption of the form $C_n^{\Sigma'}$ provided the conclusion is a rex formula χ whose associated index n_χ is lower than n. Note that this is a proper inference rule, in the sense that it is decidable whether the side conditions for the application of the rule are met.

Let us denote derivability in the enriched proof system by $\vdash_{\mathsf{coh}}$. More precisely, we write $\Phi \vdash_{\mathsf{coh}} \psi$ if for some finite $\Phi_0 \subseteq \Phi$, the sequent $\Phi_0 \vdash \psi$ is derivable in this system. First, we show that this system is sound for InqBQ.

Proposition 9.1 (Soundness)**.**
Let Σ be a (finite or infinite) signature which is function-rigid or contains identity. For all $\Phi \cup \{\psi\} \subseteq \mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$, $\Phi \vdash_{\mathsf{coh}} \psi$ implies $\Phi \models_{\mathsf{InqBQ}} \psi$

Proof. We focus on the soundness of the coherence rule (the remaining rules are discussed in Ciardelli, 2016b). We need to show that if we have $\Theta, C_n^{\Sigma'} \models_{\mathsf{InqBQ}} \chi$ for $\chi \in \mathcal{L}_{\mathsf{rex}}$, $n \geq n_\chi$, and $\Sigma' \subseteq \Sigma$, then $\Theta \models_{\mathsf{InqBQ}} \chi$. Contrapositively, suppose $\Theta \not\models_{\mathsf{InqBQ}} \chi$. By Proposition 5.3, χ is n_χ-coherent, and so by Proposition 3.5 there is a model M based on a universe W containing at most n_χ worlds, and an assignment g, such that $M \models_g \Theta$ but $M \not\models_g \psi$. Since $n \geq n_\chi$, W contains at most n worlds, and so by Proposition 6.5 we have $M \models_g C_n^{\Sigma'}$. Thus, M is a model that witnesses $\Theta, C_n^{\Sigma'} \not\models_{\mathsf{InqBQ}} \chi$. □

Next, we show that our extended proof system is also complete with respect to entailments whose conclusion is a rex formula.

Theorem 9.2 (Completeness for rex conclusions)**.**
Let Σ be a (finite or infinite) signature which is function-rigid or contains identity. For all $\Phi \subseteq \mathcal{L}_{\mathsf{InqBQ}}(\Sigma)$ and all $\chi \in \mathcal{L}_{\mathsf{rex}}(\Sigma)$, $\Phi \models_{\mathsf{InqBQ}} \chi$ implies $\Phi \vdash_{\mathrm{coh}} \chi$.

Proof. Suppose $\Phi \models_{\mathsf{InqBQ}} \chi$ and χ is a rex formula. Since χ is finitely coherent, Theorem 4.3 guarantees that $\Phi_0 \models_{\mathsf{InqBQ}} \chi$ for some finite $\Phi_0 \subseteq \Phi$. Now the set $\Phi_0 \cup \{\chi\}$ is a finite set and will thus be included in the language $\mathcal{L}_{\mathsf{InqBQ}}(\Sigma')$ for some finite $\Sigma' \subseteq \Sigma$, which we may take to include identity if Σ does. Now fix any $n \geq n_\chi$. Since $\models_{\mathsf{InqBQ}}$ is included in $\models_{\mathsf{InqBQ}_n}$, we also have $\Phi_0 \models_{\mathsf{InqBQ}_n} \chi$. Applying the completeness result for finite-bound inquisitive logics over the signature Σ' (Theorem 8.1), it follows that $\Phi_0 \vdash_n \chi$, which implies that $\Phi_0, C_n^{\Sigma'} \vdash_{\mathrm{coh}} \chi$ and thus also $\Phi, C_n^{\Sigma'} \vdash_{\mathrm{coh}} \chi$. Since χ is a rex formula with $n_\chi \leq n$ and $\Sigma' \subseteq \Sigma$, the coherence rule can be applied to obtain a proof of $\Phi \vdash_{\mathrm{coh}} \chi$. $\square$

10 Open problems

We conclude by reviewing a number of interesting questions that we have left open, and by outlining some extensions of the present work.

Recall that the coherence degree of a formula is the least cardinal κ for which the formula is κ-coherent. We have seen that for each natural number n, there are formulas of InqBQ whose coherence degree is n (for instance, the cardinality formulas $C_n^{\{R\}}$ for an arbitrary $R \in \Sigma$). There are also formulas like $\exists x Px$ that have no coherence degree, as they are not κ-coherent for any κ. One salient open question here is whether there are any formulas of InqBQ of infinite coherence degree—i.e., formulas that are κ-coherent for some infinite κ but not for any finite κ. We conjectured in Section 3 that the answer is negative, i.e., that if a formula of InqBQ is coherent at all, it is finitely coherent.

Another interesting question is whether the properties of coherence/finite coherence/n-coherence for some fixed n are (partially) decidable: is there an algorithm to (partially) decide, given a formula φ, whether it is coherent/finitely coherent/n-coherent for some fixed n?

Another open problem concerns the relation between finite coherence and the restricted existential fragment that we studied in the paper. Is the fragment expressively complete for the finitely coherent properties expressible in InqBQ? That is, is it the case that every finitely coherent formula of InqBQ is equivalent to one where the inquisitive existential occurs only within the antecedent of an implication? If not, is it possible to identify a larger syntactic fragment of InqBQ which is expressively complete in this sense?

Other questions concern the sequence of approximations InqBQ_κ for κ a cardinal. We have achieved a good understanding of the initial segment of this sequence, axiomatizing all logics InqBQ_n and showing that their intersection is strictly larger than InqBQ. On the other hand, we have said almost nothing about the transfinite part of the sequence: we know that for some κ, InqBQ_κ coincides with InqBQ, but we don't know what the least κ is for which this obtains; in particular, whether this happens already for $\kappa = \aleph_0$. A positive answer would be a significant result, analogous to the downward Löwenheim-Skolem theorem: it would mean that if an InqBQ-entailment is refutable at all, it is refutable in a countable information state. A negative answer, on the other hand, would mean that there are entailments that are invalid but can only be refuted in uncountable information states.

Another open question concerns the completeness of the proof system for InqBQ shown in Figure 1. We have shown that this system extended with the coherence rule is complete for rex conclusions. We have not shown, however, that the coherence rule is indispensable for this. This leads naturally to the following question: is the system in Figure 1 already complete for rex conclusions? If the answer is negative, this means *a fortiori* that this system is not complete for InqBQ—something which is still not known at this stage.

We conjecture that the answer is indeed negative, based on the following considerations. Consider the signature $\Sigma = \{P^{(1)}\}$ and take the formula

$$\eta \ := \ (C_2 \to C_1) \to C_1$$

where we omit reference to Σ for simplicity. This formula is valid in InqBQ. To see this, suppose an information state s does not support C_1. Then there are two worlds $w_0, w_1 \in s$ which are not duplicates. Then the state $\{w_0, w_1\} \subseteq s$ supports C_2 but not C_1, which shows that s does not support $C_2 \to C_1$. Also, η is a rex formula, since the only occurrences of an inquisitive existential are within C_2, and thus within an antecedent. We can verify that η is indeed derivable with the help of the coherence rule: we have $C_2, C_2 \to C_1 \vdash_{\text{coh}} C_1$, and since $n_{C_1} = 2$, the coherence rule allows us to discharge the assumption C_2, yielding $C_2 \to C_1 \vdash_{\text{coh}} C_1$ and thus $\vdash_{\text{coh}} (C_2 \to C_1) \to C_1$. Is η also derivable *without* the coherence rule? We strongly doubt that it is, but we do not have a proof.

Finally, it might be interesting to generalize our results by considering a *functional* notion of coherence. Given a (class-sized) function f from cardinals to cardinals, say that φ is f-coherent if for any model $M = \langle W, D, I\rangle$, the coherence condition holds for states in M with $n = f(\#D)$. This is a generalization of the notion of κ-coherence considered here, since the latter is retrieved as f_κ-coherence where f_κ is the constant κ function. On the other hand, the notion is more broadly applicable: for instance, while the formula $\exists x Px$ is not κ-coherent for any κ, it is not hard to see that it is *id*-coherent, where *id* is the identity function: that is, relative to each model M, the coherence condition is satisfied for $n = \#D$. To what extent is this functional notion of coherence informative? If we know that a formula is f-coherent for a given function f, does this allow us to draw any interesting conclusions? We leave these questions for future work.

References

Abramsky, S. and Väänänen, J. (2009). From IF to BI. *Synthese*, **167**(2), 207–230.

Ciardelli, I. (2009). Inquisitive semantics and intermediate logics. MSc Thesis, University of Amsterdam.

Ciardelli, I. (2016a). Dependency as question entailment. In S. Abramsky, J. Kontinen, J. Väänänen, and H. Vollmer, editors, *Dependence Logic: theory and applications*, pages 129–181. Springer International Publishing Switzerland.

Ciardelli, I. (2016b). *Questions in logic*. Ph.D. thesis, Institute for Logic, Language and Computation, University of Amsterdam.

Ciardelli, I. and Roelofsen, F. (2011). Inquisitive logic. *Journal of Philosophical Logic*, **40**(1), 55–94.

Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2018). *Inquisitive Semantics*. Oxford University Press.

Gabbay, D. (1981). *Semantical investigations in Heyting's intuitionistic logic*, volume 148. Springer Science & Business Media.

Galliani, P. (2012). Inclusion and exclusion dependencies in team semantics – on some logics of imperfect information. *Annals of Pure and Applied Logic*, **163**(1), 68–84.

Grädel, E. and Väänänen, J. (2013). Dependence and independence. *Studia Logica*, **101**(2), 399–410.

Grilletti, G. (2019). Disjunction and existence properties in inquisitive first-order logic. *Studia Logica*, **107**(6), 1199–1234.

Grilletti, G. (2020). *Questions & Quantification*. Ph.D. thesis, Institute for Logic, Language and Computation, University of Amsterdam.

Grilletti, G. (2021). Completeness for the classical antecedent fragment of inquisitive first-order logic. *Journal of Logic, Language, and Information*, **30**, 725–751.

Grilletti, G. and Ciardelli, I. (2021). Games and cardinalities in inquisitive first-order logic. *The Review of Symbolic Logic*, pages 1–28.

Hodges, W. (1997a). Compositional semantics for a language of imperfect information. *Logic Journal of IGPL*, **5**(4), 539–563.

Hodges, W. (1997b). Some strange quantifiers. In *Structures in logic and computer science. Lecture notes in computer science.*, pages 51–65. Springer, New York.

Kontinen, J. (2013). Coherence and computational complexity of quantifier-free dependence logic formulas. *Studia Logica*, **101**(2), 267–291.

Kontinen, J. and Väänänen, J. (2009). On definability in dependence logic. *Journal of Logic, Language and Information*, **18**(3), 317–332.

Roelofsen, F. (2013). Algebraic foundations for the semantic treatment of inquisitive content. *Synthese*, **190**(1), 79–102.

Sano, K. (2011). First-order inquisitive pair logic. In *Proceedings of the Fourth Indian Conference on Logic and its Applications*.

Troelstra, A. S. and Schwichtenberg, H. (2000). *Basic proof theory*, volume 43. Cambridge University Press.

Väänänen, J. (2007). *Dependence Logic: A New Approach to Independence Friendly Logic*. Cambridge University Press.

Yang, F. (2014). *On extensions and variants of dependence logic: A study of intuitionistic connectives in the team semantics setting*. Ph.D. thesis, University of Helsinki.

Yang, F. (2019). Negation and partial axiomatizations of dependence and independence logic revisited. *Annals of Pure and Applied Logic*, **170**(9), 1128–1149. The 23rd Workshop on Logic, Language, Information and Computation.

Yang, F. and Väänänen, J. (2016). Propositional logics of dependence. *Annals of Pure and Applied Logic*, **167**(7), 557–589.

A Cardinality formulas

In this appendix we give the details of the proof of Proposition 6.5, capturing the key property of cardinality formulas in the general case. The proof follows the same strategy as the one of Proposition 6.2, but the technical details are slightly more involved.

Proof of Proposition 6.5. We may assume without loss of generality that s contains no duplicates. We give the proof for the case in which Σ contains the identity predicate; the proof for the case in which Σ is function-rigid is analogous but simpler. To ease notation, we drop the superscript Σ.

The case $n = 0$ is obvious. For $n = 1$, we need to show that $s \models C_1 \iff \#s \leq 1$. The right-to-left direction amounts to checking that C_1 is supported at any singleton state, which is straightforward. For the left-to-right direction, suppose $s \models C_1$. We need to show that for any two worlds $w, w' \in s$, the relational structures $\mathcal{M}_w$ and $\mathcal{M}_{w'}$ coincide, which means that $w \approx w'$.

First, since $\forall x \forall y ?(x = y)$ is a conjunct of C_1, s supports this sentence. This implies that the extensions $=_w$ and $=_{w'}$ are the same, so the domains of the structures $\mathcal{M}_w$ and $\mathcal{M}_{w'}$ coincide (recall that these domains are the quotients of D relative to $=_w$ and $=_{w'}$, respectively). Moreover, for any other predicate symbol R, s supports $\forall\overline{x}?R(\overline{x})$, which implies that $R_w = R_{w'}$, and thus also the interpretations of R in the two quotients coincide. If f is a rigid function symbol, then its extension is the same at every world, so in particular $\mathsf{f}_w = \mathsf{f}_{w'}$. Finally, consider a non-rigid function symbol f. We need to show that the interpretations of f in $\mathcal{M}_w$ and in $\mathcal{M}_{w'}$ coincide. Call these interpretations $f_w^{\approx}$ and $f_{w'}^{\approx}$. For definiteness, suppose f is unary (other cases are analogous). Since the formula $\forall y\exists z(f(y) = z)$ is a conjunct of C_1, this formula is supported by s. This implies that for any $a \in D$ there is a $b \in D$ such that $\forall u \in s$ we have $f_u(a) =_u b$, and thus, in particular, $f_w(a) =_w b$ and $f_{w'}(a) =_{w'} b$. Now take an arbitrary equivalence class $[a]$ in the common domain of $\mathcal{M}_w$ and $\mathcal{M}_{w'}$. By the previous condition, there is an element $b \in D$ such that $f_w^{\approx}([a]) = [b]$ and $f_{w'}^{\approx}([a]) = [b]$, which implies $f_w^{\approx}([a]) = f_{w'}^{\approx}([a])$. Since $[a]$ was an arbitrary element of the domain, we have $f_w^{\approx} = f_{w'}^{\approx}$. This completes the proof of the fact that for all $w, w' \in s$, $\mathcal{M}_w = \mathcal{M}_{w'}$, that is, w and w' are duplicates. Since s contains no duplicate worlds, $\#s \leq 1$.

For the inductive step, assume the property holds for the formulas C_k with $k < n$ and consider the formula C_n for $n > 2$. In case $s = \emptyset$ or s contains exactly one world, we reason as in the proof of Proposition 6.2. So suppose that s contains at least two worlds and at most n. Let w_0 and w_1 be two distinct worlds in s. Since we are assuming that s contains no duplicates, $\mathcal{M}_{w_0}$ and $\mathcal{M}_{w_1}$ are distinct models. This might be the case for two distinct reasons: (i) the extensions of a relation R_{w_0} and R_{w_1} are distinct (thus the structures $\mathcal{M}_{w_0}$ and $\mathcal{M}_{w_1}$ have different domains in case R is the identity, or the extensions of R in the two structures are different) or (ii) the extensions of all predicate symbols coincide in the two worlds, but the quotients $f_{w_0}^{\approx}$ and $f_{w_1}^{\approx}$ of the interpretations a non-rigid function symbol are different.

- If case (i) applies, since $R_{w_0} \neq R_{w_1}$ we can find a tuple of elements $\overline{a}$ such that $\overline{a} \in (R_{w_0} - R_{w_1})$ or $\overline{a} \in (R_{w_1} - R_{w_0})$. Without loss of generality, suppose the former and define $s_R^+ = \{w \in s \mid \overline{a} \in R_w\}$ and $s_R^- = \{w \in s \mid \overline{a} \notin R_w\}$. Note that s_R^+ and s_R^- are both non-empty and form a partition of s. By the same reasoning as in the proof of Proposition 6.2, for $k := \#s_R^+$ we have that $s \models_{[\overline{x}\mapsto\overline{a}]} R(\overline{x}) \to C_k$ and $s \models_{[\overline{x}\mapsto\overline{a}]} \neg R(\overline{x}) \to C_{n-k}$. This implies that $s \models \exists\overline{x}\bigvee_{i=1}^{n}[(R(\overline{x}) \to C_k) \wedge (\neg R(\overline{x}) \to C_{n-k})]$, and thus $s \models C_n$.
- If case (ii) applies, we have $f_{w_0}^{\approx} \neq f_{w_1}^{\approx}$. For definiteness, suppose f is unary (the general case is analogous). We can then find an element a such that $f_{w_0}^{\approx}([a]) \neq f_{w_1}^{\approx}([a])$. Since we are assuming that the extension of identity is the same in w_0 and w_1 it follows that $f_{w_0}(a) \neq_{w_1} f_{w_1}(a)$. Thus, for $b := f_{w_0}(a)$ we have that $f_{w_0}(a) =_{w_0} b$ and $f_{w_1}(a) \neq_{w_1} b$. Now we can define the sets $s_f^+ = \{w \in s \mid f_w(a) =_w b\}$ and $s_f^- = \{w \in s \mid f_w(a) \neq_w b\}$. As in the previous case, s_f^+ and s_f^- are non-empty and for $k = \#s_f^+$ we have $s \models_{[y\mapsto a, z\mapsto b]} (f(y) = z \to C_k) \wedge (f(y) \neq z \to C_{n-k})$, which implies $s \models C_n$.

So in both cases we have that $s \models C_n$, showing that the property holds under the assumption that s has at most n worlds.

Finally, suppose that s has more than n distinct worlds. Consider an arbitrary predicate symbol R and an arbitrary tuple of elements $\overline{a}$ of size the arity of R, define the sets s_R^+ and s_R^- as above. These sets have empty intersection and $s = s_R^+ \cup s_R^-$, thus $\#s = \#s_R^+ + \#s_R^-$. In particular, for every choice of a value $k \in \{1, \ldots, n-1\}$ we have either that $\#s_R^+ > k$ or $\#s_R^- > n-k$. By the induction hypothesis, this implies that $s_R^+ \not\models C_k$ or $s_R^- \not\models C_{n-k}$. Finally since $s_R^+ \models_{[\overline{x}\mapsto\overline{a}]} R(\overline{x})$ and $s_R^- \models_{[\overline{x}\mapsto\overline{a}]} \neg R(\overline{x})$, it follows that $s \not\models_{[\overline{x}\mapsto\overline{a}]} (R(\overline{x}) \to C_k) \wedge (\neg R(\overline{x}) \to C_{n-k})$. Since $\overline{a}$ and k were arbitrary, this entails that

$$s \not\models \exists\overline{x} \bigvee_{i=1}^{n-1} [(R(\overline{x}) \to C_i) \wedge (\neg R(\overline{x}) \to C_{n-i})]$$

Consider now an arbitrary non-rigid function symbol f, an arbitrary tuple of elements $\overline{a}$ of size the arity of f and an arbitrary element b. Define the sets $s_f^+ := \{w \in s \mid f_w(\overline{a}) =_w b\}$ and $s_f^- := \{w \in s \mid f_w(\overline{a}) \neq_w b\}$. Once again, these sets have empty intersection and $s = s_f^+ \cup s_f^-$, thus reasoning in the same way as above we can conclude that

$$s \not\models \exists\overline{y}z \bigvee_{i=1}^{n-1} [(f(\overline{y}) = z \to C_i) \wedge (f(\overline{y}) \neq z \to C_{n-i})]$$

We have thus shown that s cannot support any disjunct of the formula C_n, which allows us to conclude $s \not\models C_n$ as desired. □

B Operations on state formulas

In this appendix we show explicitly that the set of state-sentences for a saturated n-theory Δ is closed under conjunction and implication, up to equivalence in Δ. Since all the formulas involved here are classical, this is essentially an exercise in classical logic about the properties of disjunctions of sentences which, relative to a background theory, are jointly exhaustive and mutually exclusive. Nothing properly inquisitive plays a role here.

Proof of Lemma 8.8. First consider the case of conjunction. If β or γ is $\bot$ then $\beta \wedge \gamma \dashv\vdash_\Delta \bot$ and $\bot \in \mathcal{L}_s^\Delta$. Otherwise, $\beta = \alpha_{i_1} \vee \cdots \vee \alpha_{i_k}$ and $\gamma = \alpha_{j_1} \vee \cdots \vee \alpha_{j_h}$. By distributivity of $\wedge$ over $\vee$ for classical formulas, which is provable in our system, we have

$$(\beta \wedge \gamma) \dashv\vdash_\Delta \bigvee_{l \leq k, g \leq h} (\alpha_{i_l} \wedge \alpha_{j_g})$$

Now consider a disjunct $\alpha_{i_l} \wedge \alpha_{j_g}$. If α_{i_l} is distinct from α_{j_g} then $\neg(\alpha_{i_l} \wedge \alpha_{j_g}) \in \Delta$, and so $(\alpha_{i_l} \wedge \alpha_{j_g}) \dashv\vdash_\Delta \bot$, which ensures that the disjunct can be removed from the disjunction while preserving equivalence in Δ. If α_{i_l} is identical to α_{j_g} then $(\alpha_{i_l} \wedge \alpha_{j_g}) \dashv\vdash_\Delta \alpha_{i_l}$, and so the disjunct can be replaced by α_{i_l} in the disjunction preserving equivalence in Δ. Let $\beta \sqcap \gamma$ be the formula obtained by removing or replacing each disjunct in this way (in case all the disjuncts end up being removed, we let $\beta \sqcap \gamma$ be $\bot$). This formula is either $\bot$ or a disjunction of world-sentences α_i, and so it is in $\mathcal{L}_s^\Delta$. Moreover, we have shown that $(\beta \wedge \gamma) \dashv\vdash_\Delta (\beta \sqcap \gamma)$.

Consider now the case of implication. If β is $\bot$, then $\beta \to \gamma \dashv\vdash_\Delta \top$, and $\top \dashv\vdash_\Delta \bigvee_{i=1}^m \alpha_i \in \mathcal{L}_s^\Delta$ by the properties of world-sentences. If γ is $\bot$, then $\beta \to \gamma \dashv\vdash_\Delta \bot$ and $\bot \in \mathcal{L}_s^\Delta$. So we can assume that both β and γ are disjunctions of world-sentences.

Consider first the case that both β and γ consist of a single disjunct, that is, $\beta = \alpha_i$ and $\gamma = \alpha_j$. We start by noticing that $\neg\alpha_i \dashv\vdash_\Delta \bigvee_{i' \neq i} \alpha_{i'}$. Indeed, since the relevant formulas are classical and our system includes a complete proof system for classical first-order logic, and using the properties of world-sentences, we have:

$$\neg\alpha_i \to \bigvee_{i' \neq i} \alpha_{i'} \;\;\dashv\vdash_\Delta\;\; \neg\neg\alpha_i \vee \bigvee_{i' \neq i} \alpha_{i'} \;\;\dashv\vdash_\Delta\;\; \bigvee_{i=1}^{m} \alpha_{i'} \;\in\; \Delta$$

$$\bigvee_{i' \neq i} \alpha_{i'} \to \neg\alpha_i \;\;\dashv\vdash_\Delta\;\; \bigwedge_{i' \neq i} (\alpha_{i'} \to \neg\alpha_i) \;\;\dashv\vdash_\Delta\;\; \bigwedge_{i' \neq i} \neg(\alpha_{i'} \wedge \alpha_i) \;\in\; \Delta$$

Given this we have that:

$$\begin{array}{rcl}
\beta \to \gamma & = & \alpha_i \to \alpha_j \\
 & \dashv\vdash_\Delta & \neg\alpha_i \vee \alpha_j \\
 & \dashv\vdash_\Delta & \left(\bigvee_{i' \neq i} \alpha_{i'}\right) \vee \alpha_j \\
 & \dashv\vdash_\Delta & \begin{cases} \bigvee_{i'=1}^{m} \alpha_{i'} & \text{if } i = j \\ \bigvee_{i' \neq i} \alpha_{i'} & \text{if } i \neq j \end{cases}
\end{array}$$

So, whether or not $i = j$, the formula $\alpha_i \to \alpha_j$ is provably equivalent to a formula in $\mathcal{L}_s^\Delta$. So for world formulas α_i, α_j we defined the formula $\alpha_i \sqsupset \alpha_j$.

As for the general case, assume that $\beta = \alpha_{i_1} \vee \cdots \vee \alpha_{i_k}$ and $\gamma = \alpha_{j_1} \vee \cdots \vee \alpha_{j_h}$. In this case we have that:

$$\begin{array}{rcl}
\beta \to \gamma & = & (\alpha_{i_1} \vee \cdots \vee \alpha_{i_k}) \to (\alpha_{j_1} \vee \cdots \vee \alpha_{j_h}) \\
 & \dashv\vdash_\Delta & \bigvee_{j'=1}^{h} \bigwedge_{i'=1}^{k} \left(\alpha_{i_{i'}} \to \alpha_{j_{j'}}\right) \\
 & \dashv\vdash_\Delta & \bigvee_{j'=1}^{h} \bigwedge_{i'=1}^{k} \left(\alpha_{i_{i'}} \sqsupset \alpha_{j_{j'}}\right) \\
 & \dashv\vdash_\Delta & \bigvee_{j'=1}^{h} \bigsqcap_{i'=1}^{k} \left(\alpha_{i_{i'}} \sqsupset \alpha_{j_{j'}}\right)
\end{array}$$

where the formulas $\alpha_{i_{i'}} \sqsupset \alpha_{j_{j'}}$ are the one previously defined. Note that the last formula is a classical disjunction of formulas in $\mathcal{L}_s^\Delta$, and thus is provably equivalent to a single formula in $\mathcal{L}_s^\Delta$. This concludes the proof. $\square$

C Saturation lemma

In this appendix we provide a proof of Lemma 8.19. This proof is an adaptation of the proof provided in Gabbay (1981, Section 3.3, Theorem 2) for the intuitionistic first-order theories with constant domains. To shorten the proofs and lighten the notation, we will omit passages regarding basic properties of the operators (e.g., associativity and commutativity of $\mathbin{\rotatebox[origin=c]{0}{$\vee$}\!\!\!\vee}$). Moreover, since the index n does not play an essential role in the proof, to lighten the notation we use the symbol $\vdash$ in place of $\vdash_n$. Finally, for the rest of the section we adopt the following notation: given Π and Ξ two sets of formulas, we indicate with $\Pi \vdash_n \Xi$ that there exists $\xi_1, \dots, \xi_k \in \Xi$ such that $\Pi \vdash_n \xi_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \xi_k$.

Before we tackle the main proof, we need the following technical result.

Lemma C.1. Consider a set of sentences $\Phi \cup \Psi \cup \{\chi\} \subseteq \mathcal{L}^A$. If $\Phi, \chi \vdash \Psi$ and $\Phi \vdash \Psi, \chi$ then $\Phi \vdash \Psi$.

Proof. By hypothesis, for some $\psi_j, \psi'_{j'} \in \Psi$, we have $\Phi, \chi \vdash \psi_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \psi_k$ and $\Phi \vdash \psi'_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \psi'_{k'} \vee\!\!\!\vee \chi$. Defining $\psi := \psi_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \psi_k$ and $\psi' := \psi'_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \psi'_{k'}$, we can abbreviate the previous expressions as $\Phi, \chi \vdash \psi$ and $\Phi \vdash \psi' \vee\!\!\!\vee \chi$ respectively. Combining these facts we get

$$\dfrac{\Phi \vdash \psi' \vee\!\!\!\vee \chi \qquad \dfrac{\dfrac{\dfrac{}{\psi' \vdash \psi'}}{\Phi, \psi' \vdash \psi'}}{\Phi, \psi' \vdash \psi \vee\!\!\!\vee \psi'} \qquad \dfrac{\Phi, \chi \vdash \psi}{\Phi, \chi \vdash \psi \vee\!\!\!\vee \psi'}}{\Phi \vdash \psi \vee\!\!\!\vee \psi'}$$

And since $\psi \vee\!\!\!\vee \psi' = \psi_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \psi_k \vee\!\!\!\vee \psi'_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \psi'_{k'}$ is a disjunction of formulas in Ψ, we obtain $\Phi \vdash \Psi$. ☐

We are now ready to prove Lemma 8.19.

Proof of Lemma 8.19. Our aim is to find a n-saturated theory Δ in the language $\mathcal{L}^A$ such that $\Phi \subseteq \Delta$ and $\psi \notin \Delta$. Fix an enumeration $B_1, B_2, \dots$ of the sentences of $\mathcal{L}^A$.[13] We define inductively a chain of pairs of theories $\langle \Delta_i, \Theta_i \rangle$ indexed by $i \in \mathbb{N}$ such that:

1. $\Delta_i \nvdash \Theta_i$.
2. For every index i, $\Delta_i \subseteq \Delta_{i+1}$ and $\Theta_i \subseteq \Theta_{i+1}$.
3. $B_i \in \Delta_{i+1} \cup \Theta_{i+1}$.

The plan is to take $\Delta := \bigcup_{i \in \mathbb{N}} \Delta_i$. During the construction we will impose some additional conditions to ensure Δ to be a n-saturated theory.

We start the construction by letting $\langle \Delta_0, \Theta_0 \rangle := \langle \Phi, \{\psi\} \rangle$. Conditions 2 and 3 are trivially satisfied. By assumption $\Phi \nvdash \psi$, and so Condition 1 is satisfied.

Suppose we already defined $\langle \Delta_m, \Theta_m \rangle$ with the properties above. By Lemma C.1, we cannot have both that $\Delta_m \vdash \Theta_m, B_m$ and that $\Delta_m, B_m \vdash \Theta_m$. So we continue the proof by considering two possible (non mutually exclusive) cases: if $\Delta_m \nvdash \Theta_m, B_m$ and if $\Delta_m, B_m \nvdash \Theta_m$.

1. *Case* $\Delta_m \nvdash \Theta_m, B_m$. We distinguish two sub-cases, depending on whether B_m is of the form $\forall x \varphi$ or not.

[13] Notice that this can be done without the use of the Axiom of Choice since we are considering a countable signature Σ and a countable set of parameters A.

(a) *Case* $B_m = \forall x\varphi$. Consider a fresh parameter $\mathsf{a} \in A$ (that is, an element not appearing in $\Delta_m \cup \Theta_m \cup \{B_m\}$) and define $\Delta_{m+1} := \Delta_m$ and $\Theta_{m+1} := \Delta_m \cup \{B_m, \varphi[\mathsf{a}/x]\}$.
Conditions 2 and 3 are respected. We want to show that also condition 1 holds, i.e., $\Delta_{m+1} \nvdash \Theta_{m+1}$. Towards a contradiction assume this is not the case, that is, for some $\theta_1, \dots, \theta_k \in \Theta_m$ and defining $\theta := \theta_1 \varovee \cdots \varovee \theta_m$ we have $\Delta_m \vdash \theta \varovee \forall x\varphi \varovee \varphi[\mathsf{a}/x]$. From this it follows:

$$\dfrac{\dfrac{\Delta_m \vdash \theta \varovee \forall x\varphi \varovee \varphi[\mathsf{a}/x]}{\Delta_m \vdash \forall x(\theta \varovee \forall x\varphi \varovee \varphi)}}{\Delta_m \vdash \theta \varovee \forall x\varphi}$$

So in particular $\Delta_m \vdash \Theta_m, B_m$, which is a contradiction. We have thus obtained $\Delta_{m+1} \nvdash \Theta_{m+1}$, which is exactly condition 1.

(b) *Case* $B_m \neq \forall x\varphi$. In this case we simply define $\Delta_{m+1} := \Delta_m$ and $\Theta_{m+1} := \Theta_m \cup \{B_m\}$. Conditions 1-3 follow by construction.

2. *Case* $\Delta_m, B_m \nvdash \Theta_m$. Once again, we distinguish two sub-cases, this time depending on whether B_m is of the form $\exists x\varphi$ or not.

(a) *Case* $B_m = \exists x\varphi$. Consider a fresh parameter $\mathsf{a} \in A$ and define $\Delta_{m+1} := \Delta_m \cup \{B_m, \varphi[\mathsf{a}/x]\}$ and $\Theta_{m+1} := \Theta_m$.
Clearly conditions 2 and 3 are respected. We want to show that also condition 1 holds, i.e., $\Delta_{m+1} \nvdash \Theta_{m+1}$. Towards a contradiction assume this is not the case, that is, for some $\theta_1, \dots, \theta_k \in \Theta_m$ and defining $\theta := \theta_1 \varovee \cdots \varovee \theta_k$ we have $\Delta_m, \exists x\varphi, \varphi[\mathsf{a}/x] \vdash \theta$. From this it follows:

$$\dfrac{\Delta_m, \exists x\varphi, \varphi[\mathsf{a}/x] \vdash \theta \qquad \dfrac{\exists x\varphi \vdash \exists x\varphi}{\Delta_m, \exists x\varphi \vdash \exists x\varphi}}{\Delta_m, \exists x\varphi \vdash \theta}$$

So in particular $\Delta_m, B_m \vdash \Theta_m$, which is a contradiction. So by contradiction we have $\Delta_{m+1} \nvdash \Theta_{m+1}$, which is exactly condition 1.

(b) *Case* $B_m \neq \exists x\varphi$. Define $\Delta_{m+1} := \Delta_m \cup \{B_m\}$ and $\Theta_{m+1} := \Theta_m$. Conditions 1-3 follow by construction.

Now let $\Delta := \bigcup_{i \in \mathbb{N}} \Delta_i$ and $\Theta := \bigcup_{i \in \mathbb{N}} \Theta_i$. Note that $\Delta \nvdash \Theta$, for otherwise there would be a finite m such that $\Delta_m \vdash \Theta_m$. We will show that Δ is an n-saturated theory with $\Phi \subseteq \Delta$ and $\psi \notin \Delta$.

First note that $\Phi = \Delta_0 \subseteq \Delta$. Since $\psi \in \Theta$ by construction and $\Delta \nvdash \Theta$, it follows that $\psi \notin \Delta$. What is left to show is that Δ is n-saturated.

By condition 3, every sentence of $\mathcal{L}^A$ is an element of $\Delta \cup \Theta$. This, together with $\Delta \nvdash \Theta$ and $\Theta \neq \emptyset$, ensures that Δ is deductively closed and $\bot \notin \Delta$.

For the disjunction property, suppose $\Delta \vdash \varphi \varovee \psi$. By contradiction, assume $\varphi, \psi \notin \Delta$, which in turn implies $\varphi, \psi \in \Theta$. In particular we would have $\Delta \vdash \Theta$, which is a contradiction; thus at least one among φ and ψ has to be in Δ. As φ, ψ were arbitrary formulas, the Disjunction property holds.

For the existence property, suppose $\exists x\varphi \in \Delta$. Suppose $B_m = \exists x\varphi$. We have $B_m \in \Delta_{m+1} \cup \Theta_{m+1}$ by condition 3. But if $B_m \in \Theta_{m+1}$ were the case, we would have $\Delta \vdash \Theta_{m+1}$ and consequently $\Delta \vdash \Theta$, which is a contradiction. So it follows that $B_m \in \Delta_{m+1}$. In particular, following the inductive construction presented above (case 2a), we have that $\Delta_{m+1} := \Delta_m \cup \{B_m, \varphi[\mathsf{a}/x]\}$ for some $\mathsf{a} \in A$. And so we have $\varphi[\mathsf{a}/x] \in \Delta_{m+1} \subseteq \Delta$. Since $\exists x\varphi$ is an arbitrary existential sentence, Δ has the existence property.

The normality condition follows from considerations analogous to the ones in the previous paragraph. $\square$